




Society for
NeuroEconomics

NEUROSCIENCE • PSYCHOLOGY • ECONOMICS



Abstract Book

21st Annual Meeting
October 13 – 15, 2023

Vancouver, British Columbia, Canada

neuroeconomics.org |  #SNE2023 | @socforneuroecon

■ Table of Contents

SPEAKERS1

ORAL SESSIONS

ORAL SESSION 13

ORAL SESSION 26

ORAL SESSION 311

ORAL SESSION 414

POSTER SESSIONS

POSTER SESSION 117

POSTER SESSION 234

POSTER SESSION 351

SYMPOSIA

SYMPOSIUM 167

SYMPOSIUM 270



Speakers

2023 Kavli Speaker

Gendered Work Assignments: Cause and Consequence

Lise Vesterlund

University of Pittsburgh

All organizations have work that no one wants to do: planning the office party, screening interns, attending to that time-consuming client, or simply helping others with their work. From office housework to important assignments that inevitably go unrewarded, a woman, most often, takes on these tasks. The talk will present research on the source and consequences of these differences, and will discuss the small, yet significant, changes that can help women and their organizations reach their potential.

Workshops

Perceptual decisions as action-perception loops in neural circuits

Hendrikje Nienborg

National Eye Institute

Studies of the neural basis of perceptual decision-making often treat each decision as a sensory-to-motor arc: the task is broken down into (independent) trials, and on each trial neural circuits extract the sensory information, compute the decision, and inform the action to report the decision. This reductionist approach is pragmatic. But it deemphasizes recurrency in the brain, flexibility to solve many tasks, and serial dependence in behavior. These phenomena play an important role during decision making in natural settings. In my talk, I will review advances in our understanding of the neural basis of perceptual decision-making, and highlight current challenges and open questions. I will argue that viewing perceptual decision-making as a recurrent process between multiple neural modules, rather than an arc, may provide a fruitful path towards addressing these challenges.

Simplicity Equivalents

Ryan Oprea

University of California, Santa Barbara

We provide evidence that the signature anomalies of prospect theory are not special phenomena of risk. They also arise (and often with equal strength) when subjects evaluate deterministic monetary payments that have been disaggregated to resemble lotteries. Thus, we find, e.g., apparent probability weighting in settings without probabilities and loss aversion in settings without loss. Across subjects, the appearance of anomalies in these deterministic tasks strongly predicts their appearance in lotteries. These findings suggest that much of the behavior described by prospect theory may be driven by the complexity of evaluating lotteries rather than by risk or risk preferences.

Oral Session 1

O1.1 - Attentional Discounting in Gains, Attentional Amplification in Losses

Brenden Eum¹, Stephen Gonzalez², Antonio Rangel¹

¹California Institute of Technology, ²Stanford University

Introduction: Simple choices between positively valued options are common in our daily lives and are susceptible to robust attentional choice biases. Recent evidence has shown that these attentional effects on choices are causal. However, we also encounter choices between negatively valued options. We investigate whether attentional choice biases are similar between choices in gains versus in losses.

Methods:

Eye-Tracking Task

Subjects complete 400 trials under two conditions (gain vs. loss; see Fig. 1). In the gain condition, two circles (100 dots each) appear on the left and right. White dots represent the probability of \$0; green dots represent the probability of \$10. Loss trials are similar, but green dots are replaced with red, representing the probability of -\$10. Subjects have free response time.

70 subjects were recruited and paid according to two random trials: 1 gain, 1 loss.

Computational Model

We use the Attentional Drift-Diffusion-Model (aDDM), where value comparison is modulated by the location of one's gaze. Subjects integrate noisy value signals into an evidence accumulator that evolves over time, E_t . Once E_t crosses one of two pre-specified boundaries, a choice is made. The process evolves according to:

$$E_t = E_{t-1} + \hat{V}_i \mu_t + \hat{\lambda}_i \mu_t$$

where $\mu_t \sim N(0, \sigma^2)$. The drift ($\hat{V}_i \mu_t$) depends on fixation location. If the subject is looking left at time t , then $\hat{V}_i \mu_t = d(V_L - \hat{V}_i, V_R)$, where d controls the speed of integration, V_i is the value of option i , and \hat{V}_i is an attentional discounting parameter. If the subject is looking right at time t , then drift is instead $\hat{V}_i \mu_t = d(\hat{V}_i - V_R)$.

Results:

Choice Biases

Consistent with previous findings, we find evidence of last fixation bias and net fixation bias in the gain trials (see Fig. 2). However, contrary to the predictions of the aDDM, we find an attentional bias away from the fixated option in losses. In other words, last and net fixation bias are still present in loss trials.

aDDM

We fit the aDDM separately to each subject by condition (see Fig. 3). Crucially, we find $\hat{\lambda}_{\text{gain}} < 1$ and $\hat{\lambda}_{\text{loss}} > 1$ for most subjects, meaning that the value of the nonfixated option is "discounted" in gains and "amplified" in losses.

Discussion:

Contrary to the pre-dictions of the aDDM, we found an attentional bias away from the fixated option in losses. That being said, choices and response times can be captured by an aDDM using a non-constant attentional discounting parameter.

It is possible that there is a fundamental difference in the role of attention in choices in gains versus in losses. Alternatively, subjects may be switching the evidence they choose to accumulate in the two conditions. We are exploring these hypotheses in subsequent work.

O1.2 - Pre-acquired functional connectivity predicts choice inconsistency

Asaf Madar¹, Vered Kurtz-David², Adam Hakim¹, Dino Levy¹, Ido Tavor¹

¹Tel Aviv University, ²New York University

Economic choice theories usually assume that humans maximize utility in their choices. However, studies have shown that humans make inconsistent choices, leading to suboptimal behavior, even without context dependent manipulations. Previous studies showed that brain activity in value and motor networks is associated with inconsistent choices at the moment of choice. Here we investigated if the neural predispositions, measured before a choice task, can predict choice inconsistency in a later risk choice task, suggesting it is an inherent trait represented in the synchrony of brain networks. Using functional connectivity (FC) measures from resting-state functional magnetic resonance imaging (rsfMRI), derived before any choice was made, we aimed to predict subjects' inconsistency levels in a later-performed choice task. We hypothesized that rsfMRI FC measures extracted from value and motor brain areas would predict inconsistency. Forty subjects completed a rsfMRI scan before performing a risky-choice task. We compared machine-learning models that were trained on FC that included only the hypothesized value and motor regions with models trained on whole-brain FC. We found significant correlations between subjects' true and predicted inconsistency levels for both model types: value-motor FC ($r = 0.363, p = 0.0349$) and whole-brain FC ($r = 0.424, p = 0.0132$). Moreover, feature importance analysis revealed that even the whole-brain model relied mostly on FC between value and motor areas. For external validation, we used an external dataset and model for the brain-behavior relationship, deploying a neural network pre-trained on FC matrices of 40,000 subjects and fine-tuning it on our data. This method also showed significant predictions of inconsistency levels ($r = 0.618, p = 0.0001$). Together, our results show that subjects' tendency for inconsistent choices can be predicted before they performed a single choice, from the predispositions of their nervous system. Lastly, we show that the synchrony between the value and motor networks plays a crucial role in this tendency.

O1.3 - A Grid-like Code for Value-Based Decision Making

Mark Orloff¹, Seongmin Park¹, Jake Blumwald¹, Philippe Domenech², Erie Boorman¹

¹University of California, Davis, ²University of Manitoba

Neuroeconomic studies propose that decisions utilize a common neural currency, where rewarding things such as juice, money, and social stimuli are encoded using a common neural value code in the brain (Levy & Glimcher 2012). Such value coding has been identified in dorsal anterior cingulate cortex (dACC; Kennerley et al., 2014), orbitofrontal cortex (Kennerley et al., 2014; Padoa-Schioppa et al., 2006, 2008), and ventromedial prefrontal cortex (vmPFC; Chib et al., 2009; Lebreton et al., 2009). Other recent work demonstrates that a grid code, originally identified for its role in representing an animal's position and enabling path integration in physical space (Hafting et al., 2005), encodes an individual's location in abstract space (i.e., a non-physical 2D space), for example in a 'bird space' of neck length and

leg length (Constantinescu et al., 2016) or a social space of popularity and competence (Park et al., 2020, 2021). These grid-like signals in an abstract 2D space have been found in entorhinal cortex (EC) and medial prefrontal cortex (mPFC). Here, we ask if humans use a grid code to efficiently represent a 2D value space and to infer decision vectors for value-based decisions. To test this hypothesis, we administer a task where participants (N = 28) are asked to make binary choices between two shapes which are drawn from two sets of shapes that vary along two continuous dimensions (such as frequency of a grid and orientation) corresponding to reward magnitude (\$) and probability (%), respectively, while undergoing fMRI scanning. First, we utilized computational modeling to calculate the utility of participant choices according to the cumulative prospect theory model (Tversky & Kahneman, 1992). We then test if blood-oxygen-level-dependent (BOLD) activity is associated with the subject-specific utility difference between the two shapes. We observe an inverse correlation of BOLD activity and utility of the chosen – unchosen option in dACC and insula (Fig a, $p < 0.005$, uncorrected) and a correlation of BOLD activity and utility of the chosen – unchosen option in vmPFC and striatum (Fig a, $p < 0.005$, uncorrected). Additionally, we observe BOLD activity in left EC (Fig b, pTFCE, small-volume = 0.051) showing a hexagonally symmetric pattern characteristic of grid cells for direct decision vectors between options' 'positions' in the 2D value space while participants make value-based decisions. These results suggest that individuals utilize a grid code when making value-based decisions.

O1.4 - Neural mechanisms underlying the prospective estimation of self-control costs

Candace Raio¹, Lewis Leone¹, Anna Konova², Paul Glimcher¹

¹New York University, ²Rutgers University

Objective: Failures of self-control continue to be a major challenge across a number of economic and health domains. Converging work across cognitive and decision neuroscience has shown that exerting control is registered as cognitively costly. We previously demonstrated that the subjective cost of self-control can be measured behaviorally using a willingness-to-pay mechanism, and further, that stress exposure increases the perceived cost of exercising control. Here, we sought to characterize the neural circuits underlying how these costs are estimated. We tested the hypothesis that prospective control cost estimates will be encoded in more anterior prefrontal regions consistent with past precommitment work [e.g. frontopolar cortex (FPC), orbitofrontal cortex (OFC)] as well as regions that have been implicated in encoding the cognitive cost of control [dorsal anterior cingulate cortex (dACC)], rather than traditional control regions that are known to actively deploy control (e.g., dlPFC). **Methods:** Healthy dieters (n=25) first rated snack foods on health, taste and temptation level in order to select a low, medium and high-tempting food for each individual. Participants then underwent fMRI scanning while completing a self-control choice task. On each trial, participants viewed a food image that varied on temptation level (low, medium, high), quantity (small, medium, large) and duration of time with the food (1-60 min). They reported trial-by-trial willingness-to-pay (from a \$10 study endowment) to avoid the food depicted on each trial. A realization phase followed the scanning session, during which one trial (bid) was randomly selected and entered into a standard economic auction procedure (BDM), which determined whether the food was successfully avoided or not. **Results:** Brain activity was modeled with a parametric modulator of (raw) bid value during the 4s decision period when participants evaluated how much to pay to avoid control. Higher bids yielded increased activation in FPC, mOFC and dACC, pointing to a central role in these brain regions in prospectively estimating the perceived cost of self-control. Activation extracted from the dorsolateral prefrontal cortex did not differ from zero during bid decisions. **Conclusions:** Our data suggest that estimating the subjective cost of exercising self-control engage a distinct neural circuit than those traditionally involved in implementing control. Acquiring a better understanding the neural basis of these cost estimates may provide potential neural targets to help improve the success of prospective self-control strategies.



Oral Session 2

O2.1 - Neural asymmetry in aligning with generous versus selfish descriptive norms in a charitable donation task

Paloma Diaz-Gutierrez¹, Christophe Boone², Harshil Vyas², Carolyn Declerck²

¹Universiteit Antwerpen, ²University of Antwerp

Study's objective: Aligning to what others do (herding behavior) is ubiquitous among animals and humans. Social alignment has repeatedly been shown to be supported by the brain's reward system (ventral striatum), presumably because attaining synchrony generates positive feelings related to social connectedness. However, aligning with selfish others often violates principles of fairness, creating an internal conflict that is unlikely to generate positive affect. Therefore, we argue that the neural basis of social alignment to generosity is fundamentally different from alignment to selfishness. We investigated this postulated neural asymmetry of aligning with generous- versus selfish descriptive norms in an incentivized fMRI charitable donation task.

Methods: Participants (N=50) decided in 120 trials how much of a 50€ endowment to donate to real charities, and how much to keep for themselves. Descriptive social norms were manipulated by presenting the average donations (generous vs. selfish, or no information) made by others to each charity.

Results: In the no information condition, participants donated on average half of their endowment. This amount significantly increased when participants were told that others had given more than half (generous descriptive norm condition), while it significantly decreased when told others gave less than half (selfish descriptive norm condition). For the fMRI data we performed ROI-analyses, which revealed that processing the selfish norm (> generous norm) increased amygdala activation, supporting the saliency of an internal conflict. Moreover, as participants aligned more with selfish norms, they showed reduced activity in the lateral prefrontal cortex, indicating attenuated cognitive control and impulsive decision-making. Conversely, as participants aligned more with generous norms, they showed increased activity in the ventral striatum. This increase occurred above and beyond increased ventral striatum activity which was also observed the more subjects donated in the baseline condition. Alignment with selfish norms, however, was not correlated with increased ventral striatum activity.

Conclusions: Taken together, this suggests that the neural mechanisms of herding and social alignment are not symmetrical. Aligning with generosity provides a "warm glow of herding" associated with "collective giving" (exceeding the well-known "warm glow of giving"). Conversely, aligning with selfishness does not, and instead might be an impulsive decision correlating with decreased cognitive control.

O2.2 - Neural predictors of interpersonal trust across cultures

Huan Wang¹, Brian Knutson¹, Jeanne Tsai¹

¹Stanford University

Increased cooperation and trusting acquaintances on the basis of prior interaction has been associated with increased activity in mesolimbic regions including the Medial PreFrontal Cortex (MPFC) and Nucleus Accumbens (NAcc; Rilling et al., 2002; Bellucci et al., 2017), which have been implicated in valuation of lasting benefits of cooperation. Trusting strangers based on no prior information has instead been associated with increased Anterior Insula (AIns) activity, and theoretically linked to aversion to betrayal (Aimone et al., 2014). Whether the neural mechanisms which promote trust vary across individuals and situations, and even in different cultures, has received less attention. To address this, we assessed brain activity using fMRI in 25 healthy European American and 27 Chinese adults as they played 72 trials of one-shot Trust Game. On each trial, subjects saw picture of a partner face (4s), followed by information about the previous reciprocation history of that partner (i.e., the percentage of previous players who rated the partner as trustworthy; 4s), and finally a prompt to indicate how much they chose to invest in the partner (out of a \$6 endowment; 4s), which served as a proxy for trust. As predicted, European American chose to invest more often than Chinese (i.e., a high trust default; $F=20.10$, $p<.001$). Neurally, we harvested trial-to-trial variations from the three regions of interest (i.e., NAcc, AIns, and MPFC) immediately before participants made investment choices and applied multivariate mixed-effect logistic regression to predict choices on each trial. Analyses revealed a culture by AIns interaction ($t=-2.30$, $p=.022$), in which AIns activity negatively predicted choices to invest for European Americans ($B = -.56$, $p = .007$). Analyses also revealed culture by MPFC ($t=-2.18$, $p=.029$) and culture by NAcc interactions ($t=-2.93$, $p=.003$), in which activity in these regions positively predicted choices to invest for Chinese (MPFC: $B=.32$, $p < .001$; NAcc: $B=.36$, $p=.047$). These results suggest that cultural values can shape the default level of trust, which may then interact with partner's reciprocation history to violate the default. On the one hand, European Americans with high trust default distrusted bad partners, which was predicted by higher AIns activity. On the other hand, Chinese with low trust default trusted good partners, which was predicted by increased MPFC and NAcc activity. Together, these findings highlight that motivations to trust may depend on cultural defaults in trust, which may have implications for improving cross-cultural communication and cooperation.

O2.3 - The neural mechanism of social preference changes across joint versus separate evaluation

Yinmei Ni¹, Jian Li¹

¹Peking University

Social preference, a positive or negative concern for the welfare of others, has been a central concern in explaining people's social behavior. Both separate (one option is presented and evaluated at a time) and joint evaluations (two options are evaluated simultaneously) have been widely used in the study of social preference. However, whether the two evaluations reveal different social preferences, and what the neural process under these differences remains to be uncovered. To answer these questions, we recruited 90 adult college students across two experiments. In both experiments (experiment 1: behavioral task, 34 subjects; experiment 2: fMRI task, 56 subjects), subjects both rated their satisfaction with one reward option (SE) and made binary reward choices (JE), each option includes a reward for the subject and a reward for their anonymous co-player (Figure 1).

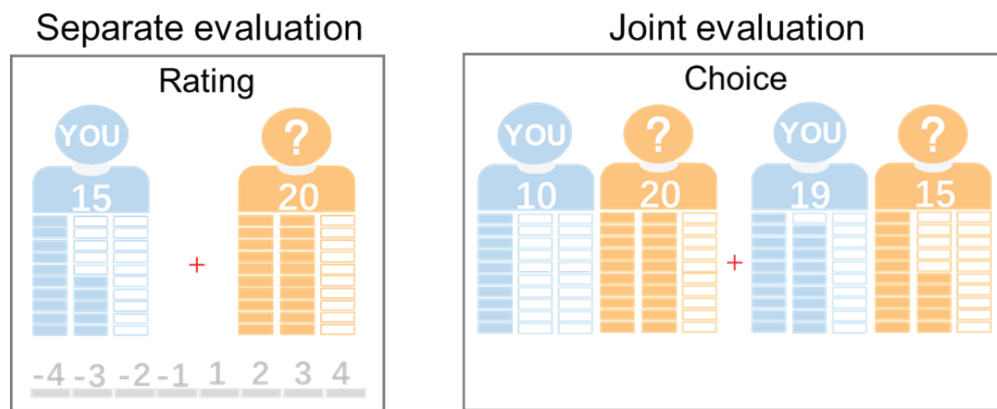


Figure 1. Experiment design

Firstly, we used a utility model () to quantify the social preferences, where α and β are the rewards for subjects and their co-player, α represents the relative weight of β which quantifies subjects' social preference. In both experiments, we found that the social preference was significantly negative in SE and significantly positive in JE (Figure 2), indicating that subjects were pro-self in SE while prosocial in JE.

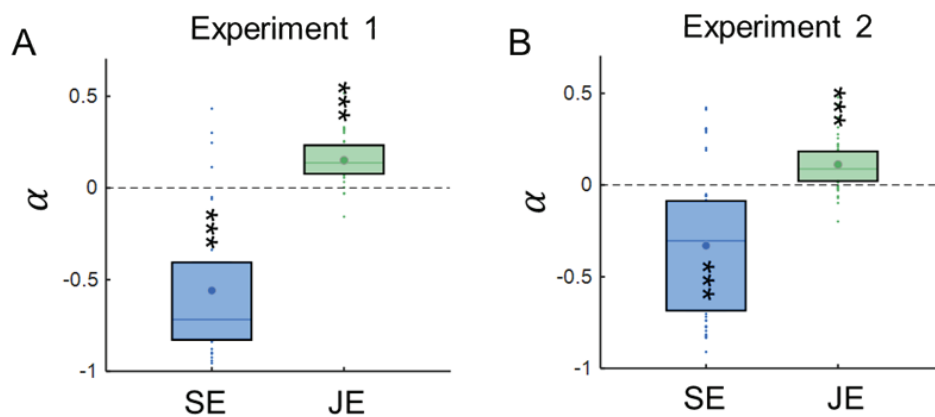


Figure 2. Social preference

Secondly, the GLM analysis revealed that, compared to SE, the encoding strength of reward for participants (Self) and reward for co-players (Other) were significantly lower in DMPFC, DLPFC and Insula (blue regions), and significantly higher in VMPFC (red region) in JE.

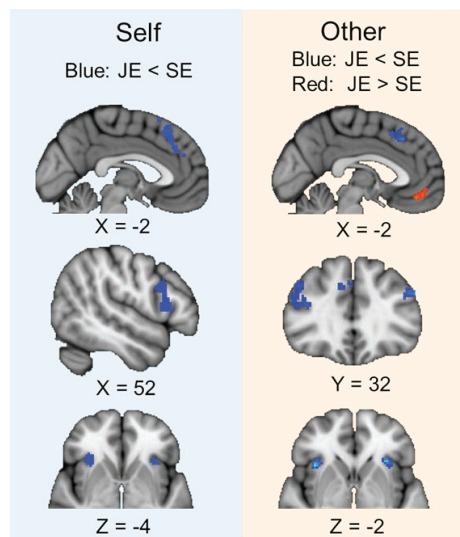


Figure 3. Neural encoding difference of Self and Other across SE and JE

Thirdly, PPI analysis revealed that the encoding difference of Other in VMPFC was modulated by the functional connectivity between rIFG /rINS and VMPFC, and the activity difference in rIFG and rINS was further modulated by their functional connectivity with rTPJ (Figure 4).




Figure 4. PPI results

Lastly, in SE and JE, utilities were encoded in separated VMPFC subregions, and the relative encoding of Self and Other in these regions was correlated with the social preference (Figure 5).

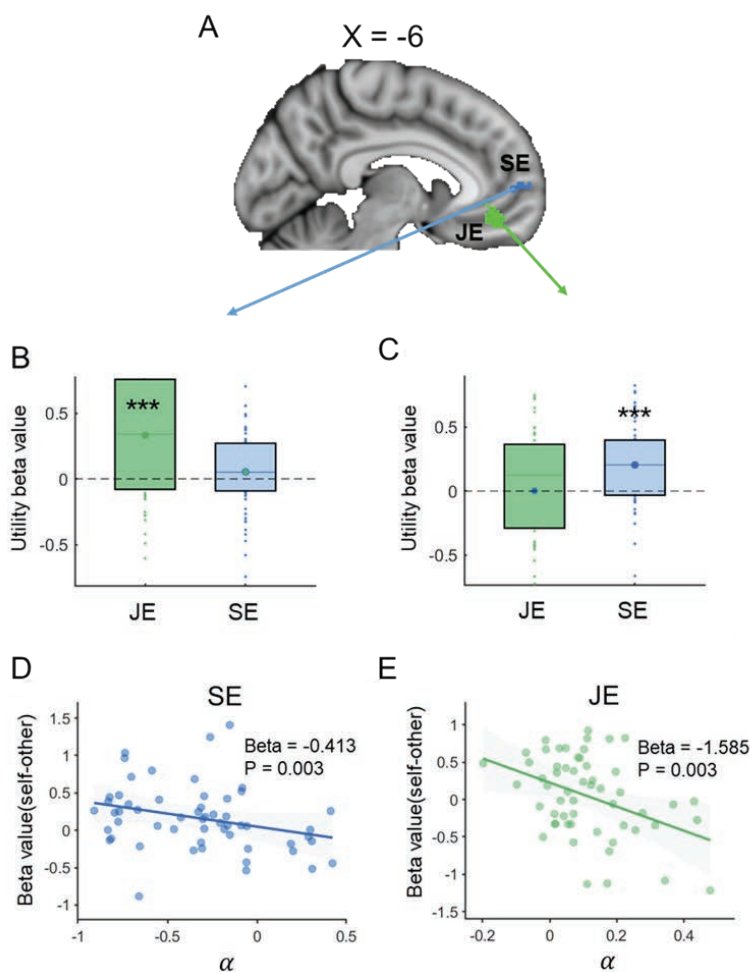


Figure 5. Utility integration in VMPFC subregions

To summarize, combining computational modeling and neuroimaging, we found that social preferences revealed by SE and JE were different, subjects were more prosocial in JE than SE, and this difference may be caused by stronger interplays across social cognition, cognitive control, and valuation brain regions in JE.

O2.4 - Stable patterns of ongoing self-generated thoughts predict variance in generosity across people and contexts

Lisa Bas¹, Ruien Wang¹, Yijun Xu¹, Anita Tusche¹

¹Queen's University

Background: People differ in their willingness to help. Why? Here, we examined the role of self-generated thoughts in explaining variance in social behaviors. Differences in ongoing thoughts (mind-wandering) have been linked to mental health, well-being, perceived loneliness, and age differences. Our study tested if stable patterns of ongoing self-generated thought predict variance in altruistic behavior across people and contexts.

Methods: In two studies, 323 healthy volunteers performed an altruistic choice task (modified dictator game, online study 1: n=166, in-person study 2: n=157). Throughout the task, participants repeatedly completed multidimensional experience sampling to assess their ongoing thoughts on 14 scales. Principal component analysis (PCA) identified common thought patterns. A linear mixed model then tested whether the prevalence of each thought pattern varied depending on the level of generosity. Moreover, we tested for the impact of choice context on thought patterns in two ways: First, we examined if online vs. in-person task environments altered the prevalence of each thought pattern. Second, in different blocks of the altruism task, people interacted with partners identified via short videos of faces that were 'identifiable sad', 'identifiable neutral', or 'unidentifiable scrambled' (control). We tested if changes in this social context modified people's thoughts.

Results: We identified four principal patterns of ongoing thoughts (PCA). Notably, the occurrence of two of these thought patterns predicted inter-individual differences in social behaviors in the altruistic choice task (generosity). For example, individuals that showed more Pattern 4 thoughts were more generous (thoughts identified as more distant other-oriented, less self-oriented, more negative, and more past-related) ($p < 0.001$). Importantly, the dominance of this generosity-related thought pattern was stable across online vs. in-person task contexts ($p = 0.92$) but significantly altered by social cues in the altruistic choice task (partner: identifiable sad, identifiable neutral, unidentifiable control) ($p < 0.001$). Seeing social facial cues that identified others in the altruism task increased the manifestation of this prosocial thought pattern. Crucially, using data from an independent free-viewing task confirmed that the identified thought patterns generalize to settings that don't require social choices, suggesting their generalizability across time and different task contexts.

Conclusions: Our results provide the first evidence for the predictive power of ongoing thought patterns to explain variance in social decision-making across people and choice settings.



Oral Session 3

O3.1 - Ambiguity aversion arises from nonlinear forward sampling of future reward states

Kenway Louie¹

¹New York University

Empirical decision-making in biological choosers depends markedly on outcome uncertainty. Such uncertainty can differ in the degree of knowledge held by a chooser: in decisions under risk, outcomes are probabilistic but those probabilities are known; in decisions under ambiguity, the probabilities themselves are unknown or uncertain. While human choosers generally exhibit aversion to both risk and ambiguity, these two types of uncertainty preferences differ in fundamental ways. Risk aversion can be rationally derived as a normative tradeoff between magnitude and probability according to individual chooser utility functions. In contrast, ambiguity aversion is normatively irrational and cannot be explained by existing models of valuation and choice.

Here, we show that **ambiguity aversion arises naturally in agents employing (1) a normalized value representation and (2) a distributional sampling of possible outcome states under ambiguity**. This nonlinear forward sampling model replicates known characteristics of empirical ambiguity aversion, including: preference for risky over ambiguous options, a quasi-linear relationship between valuation and ambiguity, and a dissociation between risk and ambiguity preferences. At the behavioral level, this model makes the counterintuitive (and testable) prediction that the degree of ambiguity aversion varies inversely with the density of forward sampling. At the neural level, the state sampling inherent in the model suggests a mechanism for ambiguity preference via distributional reinforcement learning, a process recently linked to normalized value representations. Together, these results offer a simple explanation for ambiguity preferences based on normalized value coding and forward inference, and argue for an incorporation of biologically valid value functions in computational models of decision-making.

O3.2 - Adaptive generalization and efficient learning under uncertainty

Jiwon Park¹, Dongil Chung¹

¹Ulsan National Institute of Science and Technology

People often use recognizable (or familiar) features to infer the value of novel consumables. This “generalization” strategy is known to be beneficial in stable environments, such that individuals can use previously learned rules and values in efficiently exploring new situations. However, it remains unclear whether and how individuals adjust their generalization strategy in volatile environments where previously learned information becomes obsolete. We hypothesized that individuals adaptively use generalization by continuously updating their beliefs about the credibility of the feature-defined value model at each state. To test this hypothesis, we developed a novel probabilistic learning task where participants were asked to learn about the reward contingencies assigned to a set of stimuli. Critically, each stimulus was generated as a Gabor patch that has a unique orientation between 0 and 90 degrees, and to covertly inform a monotonic (or non-monotonic) association between the stimuli orientations and the assigned reward probabilities, participants were initially trained on three anchoring pairs ([0°-0%], [45°-50%], and [90°-100%]; or [0°-50%], [45°-0%], and [90°-100%]). Subsequently, participants had to learn about four new stimuli and were informed that reward contingencies could change during the task. Our data showed that participants performed better when the novel environment remained consistent with the previously learned monotonic association between feature and reward, suggesting efficient utilization of prior knowledge ($t(83) = -3.48$, $p < 0.001$). Against other accounts, we found that individuals incorporated an arbitration mechanism between feature-based reward generalization and model-based learning based on volatility tracking. Notably, our suggested model captured differential impacts of generalization dependent on the context-volatility, such that individuals who were biased the most toward generalization showed the lowest learning errors in a stable environment ($r = -0.31$, $p = 0.053$), but the highest errors in a volatile environment ($r = 0.69$, $p < 0.001$). This work provides novel insights into the adaptive usage of generalization, orchestrating two distinctive learning mechanisms through monitoring their credibility, and highlights the potential adverse effects of over-generalization in volatile contexts.

O3.3 - Representations of information value in mouse orbitofrontal cortex during information seeking

Jennifer Bussell¹, Ethan Bromberg-Martin², Ryan Badman³, Kanaka Rajan³, Larry Abbott¹, Richard Axel⁴

¹Columbia University, ²Washington University School of Medicine, ³Icahn School of Medicine, ⁴Columbia University & Howard Hughes Medical Institute

Humans and animals seek information, even if it does not lead to increased reward and comes at an economic cost, and they therefore behave as if information is intrinsically valuable. We have investigated how the brain represents the value of information using a mouse model of information seeking. We developed an odor-based decision making task in which mice choose to receive probabilistic reward in either of two nosepoke ports, which differ only in whether they provide information via odor that reveals the reward amount but cannot be used to increase reward. Mice are cued at trial start by odor in a third port which reward ports are available, such that they learn an association between odor and available information. Mice strongly prefer the information port (67% mean preference, $p < 0.01$, $N=14$). Moreover, mice exchange water to pay for information, preferring the information port even when it offers less water reward. We fit reinforcement learning parameters to trial-level choices to model the learning of information value and found that the modeled subjective value of information matched that determined experimentally (mean 38% water value exchanged for information, $N=4$). Thus, like humans and other animals, mice display an economically sub-optimal information bias in their decision making, behaving as if information has value unto itself. We performed calcium imaging with miniature microscopes to record the activity of excitatory neurons in orbitofrontal cortex in mice performing the information seeking task. Decoding analysis and dimensionality reduction revealed a population representation of information expectation that emerged with learning (17% of all OFC excitatory neurons encode information prediction). The difference between neural activity in anticipation of information vs. no information scaled with the time animals spent in a state of uncertainty, consistent with the information representation being a scalar value signal. In addition, we also observed a representation upon the receipt of information in response to reward outcome-associated odors. There was significant similarity between the population response to positive and negative information-providing odors ($\text{corr}=0.73$, $p<0.0001$), suggesting that information is an unconditioned stimulus that reinforces the representation of information-predicting odor. Moreover, the representation predictive of information was distinct from the representation predictive of water reward ($\text{corr}=0.15$, $p<0.0001$). Thus, mice seek information of no extrinsic value, and the intrinsic value of information is encoded in a distinct representation in orbitofrontal cortex during decision making.

O3.4 - An Interplay of Nucleus Accumbens and Anterior Insula Activity Predicts Risky Choice

Leili Mortazavi¹, Charlene Wu¹, Elnaz Ghasemi¹, Brian Knutson¹

¹Stanford University

The replicability and generalizability of neuroeconomic findings has recently come under scrutiny. If neural activity that immediately precedes choice action contains information about the integrated assessment of risk, however, it should predict choice across samples, tasks, and analyses. Across five Functional Magnetic Resonance Imaging (fMRI) samples (total N=230), using two gambling tasks (Leong et al., 2016; Tom et al., 2007), and four analysis methods, we optimized a predictive model of trial-by-trial risky choice first in an original dataset (n=75), then validated it in an independent sample (n=32), and finally generalized it to a different task with various gamble properties (n=123; openly available raw data from Tom et al., 2007 & Botvinik-Nezer et al., 2020).

Convergently across datasets, we found that while a swath of overlapping regions bidirectionally respond to gain and loss outcomes of gambles, specific regions selectively predict risk-seeking (Nucleus Accumbens; NAcc) versus risk-avoidant choices (Anterior Insula; AIns) in opposite directions. We further show that the neural predictors of choice are dissociable from sensorimotor correlates of option assessment and choice execution. Further corroborating the interplay of this approach-avoidance signal in the NAcc and AIns, the ratio of NAcc: AIns pre-choice activity is higher on risk-seeking choices suggesting that the relative activity of these regions best predict choice.

Consistent with the negative main effect of AIns on risky choice, for generally risk-seeking individuals, AIns activity was high before risky choices but low before safe choices. Interestingly, more risk-averse individuals consistently showed heightened AIns activation before risky and safe choices. Nevertheless, the ratio of NAcc: AIns still predicted risk-seeking choices across individuals with divergent preferences.

These findings elucidate the nuances of how AIns activity contributes to risky choice and suggest this region may exert its influence on choice by modulating NAcc activity. By demonstrating convergent evidence across univariate and multivariate analyses, samples, and tasks, this work helps resolve a persistent neuroeconomic theoretical question by indicating that two opposing anticipatory affective signals drive risk seeking versus risk avoidance, and have implications for better predicting risky choice in future research.



Oral Session 4

O4.1 - The Dynamic Nature of Procrastination

Peiyuan Zhang¹, Wei Ji Ma¹, Yijun Lin², Falk Lieder³

¹New York University, ²University of Florida, ³Max Planck Institute for Intelligent Systems

Humans procrastinate. Procrastinators usually make little progress at the start (start late) and a significant increase in progress shortly before the deadline (rush to complete). Yet, the cognitive mechanisms underlying this time course of progress remain poorly understood. To investigate this, we developed a computational model that considers the time course of progress as the output of sequential decision-making: whether to work now (and, if so, how much) or later. If they decide to work now, they pay the cost of investing mental effort immediately but also make progress, and more work leads to more progress. If they decide not to work, they make no progress and also pay no effort cost. The amount of work on each day is derived from the Bellman equation (Bellman, 1957), which assumes that a person's goal on each day is to maximize the discounted value gained by making progress while minimizing the immediate effort cost. Our model predicts that 1) stronger temporal discounters procrastinate more, and 2) offering immediate rewards versus delayed rewards reduces procrastination (Fig 1). We confirmed our first model prediction with a real-world behavioral study (Study 1, N=93), where we observed a positive correlation between individual discount rate estimated from an inter-temporal choice task and the level of procrastination in a research participation task assigned by a psychology course ($r=.28$, $p=0.009$; Fig 2). To test our second model prediction, we created a novel experimental paradigm where people work on a self-paced, week-long online task consisting of numerous units of work that take about 3 hours to complete, and we manipulated reward timing as immediate reward upon task completion versus delayed reward (Study 2, N=611; Fig 3). We quantified the level of procrastination using three summary statistics of the time course of progress (Fig 4A): task starting day, task completion day, and mean unit completion day. Our results revealed that offering an immediate reward upon task completion helped people start the task earlier (quasi-Poisson regression: $\hat{\beta}_2(1)=6.05$, $p=0.014$). Moreover, it helps people who generally procrastinate (measured by General Procrastination Scale (Lay, 1986)) complete the task earlier (quasi-Poisson regression: $\hat{\beta}_2(1)=6.94$, $p=0.008$) and complete units of work earlier ($F(1,281)=8.38$, $p=0.004$; Fig 4B). These results support the second model prediction and also highlight how immediate rewards affect task initiation and completion differently. Taken together, our findings enhance our understanding of the dynamic nature of procrastination and offer implications for reducing procrastination.

O4.2 - Social media users' posting behavior can be as-if addictive

Felix Nitsch¹, Klaus Wertenbroch¹, Hilke Plassmann¹

¹INSEAD Business School

Social media provide various benefits such seamlessly interacting with others, expressing oneself, and learning about the latest trends. Recently, however, the dark side of social media has attracted considerable attention in public discourse: conditions and instances when social media are harmful to their users. In the current research project, we study a specific facet of social media use: creating and sharing content online (i.e., posting). The current research project documents that user posting behavior on Twitter and Instagram on average exhibits regularities reminiscent of behavioral addictions.

We web-scraped over 400.000 Instagram (IG) and 600.000 Twitter (TW) posts, following a sample 35.491 IG and 2.662 TW users over a time frame of up to about 5 (IG) or 7 (TW) years and extracted the latencies between subsequent posts in the resolution of seconds. Additionally, we extracted the most salient feature of social feedback (i.e., likes) per post and the general user characteristics.

We modelled this data using a recurrent event framework adopted from studies of disease remission and addiction relapse, identifying predictors of not only whether a user makes the next post but also how soon. Within this framework we tested 5 different hypotheses, predicting that the hazard of posting:

1. is proportional to expected likes,
2. is anti-proportional to variance of likes (uncertainty aversion),
3. dynamically adapts according to prediction errors (reinforcement learning),
4. increases with continued use, outcome-independently, and/or
5. increases when receiving less likes than expected, expressed via the hazard ratio (Hr).

We found significant evidence for H4 (IG: Hr=1.143, $p<.001$; TW: Hr=1.002, $p<.001$) and H5 (IG: Hr=1.314, $p<.001$; TW: Hr=1.432, $p<.001$; mutually exclusive to H3). Contradicting H1 and H2, posting hazard was anti-proportional to expected likes (IG: Hr=0.9996, $p<.001$; TW: Hr=0.9997, $p<.001$) and proportional to variance of likes (IG: Hr>1, $p<.001$; TW: Hr>1, $p=.680$). A combination model of H4 and H5 provided the best model fit according to the Bayesian Information Criterion. Lastly, we also found that the odds of receiving less likes than expected increase with continued use (IG: Z=53.03, $p<.001$; TW: Z=35.96, $p<.001$).

Our results suggest that social media users on average post increasingly fast and frequent independent of feedback. Further, receiving less likes than expected additionally increases posting hazard. As continued use also raises the odds of disappointed expectations, this potentially poses a vicious cycle reminiscent of behavioral addictions (i.e., "chasing the first high").

O4.3 - Neural mechanisms underlying the persuasion by healthy eating calls

Ioannis Ntoumanis¹, Alina Davydova¹, Yulia Sheronova¹, Vladimir Kosonogov¹, Ksenia Panidi¹, Anna Shestakova¹, Iiro Jääskeläinen¹, Vasily Kliucharev²

¹HSE University, ²Higher School of Economics - National Research University

Objective: As obesity rates are rising, mainly due to sugar consumption (Yu et al., 2022), it is important to understand what can influence consumers' perception of sugar. Healthy eating calls can effectively modulate individuals' willingness to pay (WTP) for sugar-containing food (Ntoumanis et al., 2022) and this modulation has been associated with the synchrony of neural responses to the healthy eating call in question (Ntoumanis et al., 2023). However, it remains unclear which exact brain areas are involved in the persuasion by naturalistic healthy eating calls.

Methods: In this experiment, 50 healthy participants performed two blocks of a bidding task, in which they had to bid on 30 sugar-containing, 30 sugar-free and 30 non-edible products, while their brain activity was recorded using functional magnetic resonance imaging (fMRI). In-between the two blocks, they listened to a 7-minute healthy eating call, by a nutritionist, emphasizing the risks of sugar consumption. The behavioral effect of the healthy eating call was quantified with the ΔWTP , i.e., the WTP for each product in the

second block of the bidding task subtracted by the WTP for the corresponding product in the first block of the bidding task. Based on the median Δ WTP for sugar-free products, participants were evenly divided into two groups (influenced/not influenced), in order to identify brain areas that were differentially synchronized during listening.

Results: At the behavioral level, participants' Δ WTP was higher for sugar-free products (Wilcoxon test statistic $W=256$, $p<.001$, effect size $r=.521$) and lower for sugar-containing products ($W=336$, $p=.016$, $r=.347$), compared to the Δ WTP for non-edible products. Moreover, using inter-subject correlation (ISC) analysis, we found that activity in higher-order brain regions was more synchronized among participants who were not influenced by the healthy eating call, compared to those who were influenced, especially in the precuneus (ISC difference $r=.033$, $p=.003$ after correction for multiple comparisons).

Conclusions: Previous studies have reported that, perceiving a narrative as fictitious activates precuneus more strongly than perceiving the same narrative as factual (Altmann et al., 2012), and that precuneus is particularly synchronized when listening to a narrative that elicits negative emotions (Jääskeläinen et al., 2020). Hence, our results indicate that the participants who were not influenced by the healthy eating call may have been suspicious of the arguments expressed in it. Overall, our study provides a promising tool to design and assess health-related advertisements before they are released to the public.

O4.4 - Why do storytelling ads persuade consumers? Evidence from brain responses to video commercials

Leo Van Brussel¹, Maarten Boksem¹, Ale Smidts²

¹Erasmus University Rotterdam, ²Erasmus University

High costs demand marketers to know whether video advertisements will have the desired effect on consumers. It has been shown that storytelling ads hook viewers, leading to more positive attitudes. Here, we use fMRI to study the mechanisms underlying the effect of narratives on video ad liking. Measuring brain responses of consumers watching video ads allows us to investigate how consumers respond to narratives, without breaking the spell of the story.

We report an analysis of two independent fMRI datasets, in which a total of 85 participants viewed and rated a total of 55 video ads. Ad liking was quantified by means of participants' star ratings (0-5 stars, half-star increments). Independent coders used a six-item narrative structure coding scale (Escalas et al., 2004) to quantify the degree to which ads contained a narrative. We used parametric whole-brain analyses to identify brain regions that activated when ads (1) were liked better by the participants (higher star rating) and (2) were more story-like (higher narrative structure score). We then used multilevel mediation analysis (Wager et al., 2009) to assess whether and where in the brain activation mediated the effect of narrative structure score on ad liking. Finally, we employed Neurosynth (Yarkoni et al., 2011) to reveal which psychological processes were most associated with our results.

A mixed-effects model confirmed that ad liking positively related to narrative structure ($b = 0.282$, $SE = 0.073$, $p < 0.001$). Whole-brain parametric analysis indexing ad liking revealed activation in Superior Temporal Gyrus (STG), Precuneus, Temporal Parietal Junction (TPJ), Temporal Pole and Amygdala for both datasets ($p < 0.001$, uncorrected; whole-brain correlation between datasets: $r = .459$). Similarly, indexing narrative structure score we found activation in STG, Precuneus and TPJ ($p < 0.001$ uncorrected; whole-brain correlation between datasets: $r = .611$). We then performed multilevel mediation and used a whole-brain search at $q < 0.05$ (FDR corrected; corresponding to $p < 0.006$) to reveal mediating activation in STG, TPJ and Cerebellum. Neurosynth decoding indicated that this activation was most strongly associated with 'theory of mind' ($r = 0.395$), and similar terms such as 'mental states' ($r = 0.327$), 'mentalizing' ($r = 0.274$) and comprehension ($r = 0.282$).

Our results reveal mentalizing as key mechanism explaining how ad liking is driven by narrative structure. This aligns with research suggesting that narrative persuasion relies on the extent to which narratives mentally transport story receivers into the world of the narrative, leading to more positive attitudes.



Poster Session 1

P1-B-1 - Individual Differences in Autistic Traits Affect Visual Engagement with Advertisements

Andrea Bazzani¹, Davide Benedetti², Paolo Frumento², Liliana Dell'osso², Giuseppe Turchetti¹, Ugo Faraguna²

¹Sant'Anna School of Advanced Studies, ²University of Pisa

Study's objective: The conceptual framework behind this study relies on differences in attentional patterns within autism spectrum. Our main goal is to investigate how individuals with different autistic traits engage with commercial advertisements and communication messages. To this purpose, we analysed three specific gaze metrics: duration of fixations, number of fixations, and fixation path (measured as the cumulative distance covered by the gaze on a stimulus or area of interest - AOI). By doing so, we aim to shed light on the possible link between individual differences in autistic traits and visual engagement towards commercial ads.

Methods: We recruited 39 participants (25.92 ± 5.63 ; %Females = 51.28), who were asked to watch a popular Italian commercial ad humorously depicting a scenario where a famous Italian architect discusses restyling Paradise with a Saint, all while enjoying a coffee mug of a well-known Italian brand. In our analysis, we focused on AOIs corresponding to faces and non-facial visual elements (e.g., hands and brand/product images).

Results: We observed a positive correlation between autistic traits and fixation time (duration) on faces ($r = 0.343$, $p = 0.035$). Participants with higher autistic traits tended to spend more time fixating on faces. Conversely, we did not observe a clear or significant relationship between autistic traits and gaze metrics for non-facial AOIs. We only reported a significant negative correlation between autistic traits and the number of fixations on brand elements ($r = -0.386$, $p = 0.017$).

To further explore differences in gaze behaviour, we calculated the ratio and the difference of gaze metrics between faces and non-faces AOIs, as well as between faces and brand elements. Participants with higher autistic traits exhibited a higher number of fixations on faces relative to brand elements ($r = 0.437$, $p = 0.006$). Also, we reported a positive significant correlation between autistic traits and the difference in fixations duration between non-faces and faces AOIs ($r = 0.337$, $p = 0.038$), as well as between faces and brand elements ($r = 0.333$, $p = 0.041$).

Conclusions: Our findings show a correlation between individual differences in autistic traits and specific gaze metrics related to engagement with commercial advertisements and communication messages. This highlights the potential of using neurophysiological measures and spectrum traits of psychiatric constructs as powerful tools for segmenting the target audience and designing personalized communication strategies accordingly.

P1-B-2 - Source Memory is More Accurate for Opinions Than for Facts

Daniel Mirny^{1,2}, Stephen Spiller²

¹IESE Business School, ²University of California, Los Angeles

Effective communication relies on consumers remembering, sharing, and applying relevant information. Source memory, the ability to link a claim to its original source, is an essential aspect of accurate recall, attitude formation, and decision making. We propose that claim objectivity, whether a claim is a fact or an opinion, affects memory for the claim's source.

This proposal follows a two-step process: (i) Correctly recalling the original source of a given claim relies on the strength of the link formed between the source and the claim during the initial encoding of information (Greene et al. 2021; Mitchell and MacPherson 2017; Pham and Johar 1997); (ii) Opinions provide more information about a source than do facts (Heiphetz et al. 2014).

Across twelve pre-registered experiments ($N=7,008$) and a variety of consumer domains, we investigate the effect of claim objectivity on source memory. The design used in each experiment was drawn from the source memory literature (e.g., Kassam et al. 2009). Stimuli were pre-tested to ensure that differences in claim objectivity were not confounded by differences in claim emotionality, valence, or arousal.

In five experiments, we find that source memory is more accurate for opinions than for facts, with no consistent effect on claim recognition memory (Exp 1 ($N=399$): $b=12.47$, $t(321)=10.54$; Exp 2a ($N=501$): $b=2.63$, $t(412)=2.23$; Exp 2b ($N=504$): $b=3.79$, $t(440)=3.43$; Exp. 2c ($N=501$): $b=3.51$, $t(464)=3.40$; Exp. 3 ($N=606$): $b=5.33$, $t(560)=3.68$).

We find support for the proposed process by manipulating facts to be more informative about sources, increasing source memory accuracy (Exp 4 ($N=403$): $b=3.99$, $t(305)=3.49$) and opinions to be less informative about sources, decreasing source memory accuracy (Exp 5 ($N=1,213$): $b=1.00$, $t(1076)=1.05$).

When forming inferences and seeking advice from sources, participants integrate new evidence more accurately based on sources' previously shared opinions than facts (Exp 6a ($N=640$): $b=4.07$, $t(590)=4.00$; Exp 6b ($N=639$): $b=2.78$, $t(589)=2.61$).

In three supplemental experiments, substantial changes to the experimental design resulted in attenuations of the main effect when using complex stimuli (Exp S1 ($N=499$): $b=-.67$, $t(321)=-.96$), when testing source memory using a cued rather than a full recall task (Exp S2 ($N=501$): $b=.03$, $t(401)=.03$), and when sources were periodical publications rather than humans (Exp S3 ($N=601$): $b=-.89$, $t(551)=-.85$).

Our results indicate that opinions are more likely to be accurately attributed to original sources than are facts.

P1-B-3 - The effect of commercials on the neural signal of value

Inbal Gur Arie¹, Adam Hakim¹, Valentina Parada, Dino Levy¹

¹Tel Aviv University

A basic aim of marketing research is to predict consumers' preferences and the success of marketing campaigns at the population level. However, traditional marketing tools have various limitations, calling for novel measures to improve predictive power. In a previous study in the lab, we developed a deep learning model that predicted subject-specific preferences of pictures of consumer products based on their EEG data. In the current study, we extend this method to video commercials. There were three parts to the current experiment. In the first part, subjects viewed a series of different consumer products and indicated their willingness to pay (WTP) for each product using a classic BDM task. Afterward, in the second part, subjects viewed commercials for these same products, and indicated, again, their WTP for each of the products and ranked how much they liked the commercial itself. During these first two parts of the experiment, we recorded subjects' brain activity using EEG. A week later, in the third part of the experiment, subjects again indicated their willingness to pay for the products, but without EEG recordings, outside the lab, as an online experiment. Our 120

subjects' findings indicate that watching commercials increased subjects' WTP, but only immediately. After a week, the subjects' WTP returned to baseline. Furthermore, commercials' liking and products' WTP are significantly but only weakly correlated ($R^2 = 0.098$, $p < 0.001$), indicating they represent different aspects of commercial preference. Moreover, we found that traditional EEG frequency bands are significantly, but weakly, correlated with subjects' WTP measured while watching the products' pictures ($R^2 = 0.017$, $p < 0.001$) or while watching their commercials ($R^2 = 0.027$, $p < 0.001$). When we conducted a logistic regression using the EEG frequency bands to predict subjects' WTP we could not predict successfully over chance levels. However, using our deep learning network, we were able to predict subjects' preferences with an average accuracy of 0.69. Moreover, we found that the low-level visual features of the commercials, such as length, color, and quality, were correlated with the subjects' WTP for the products. To conclude, people's preferences seem to generally be influenced by relevant stimuli, as commercials, while their liking and product value do not necessarily go hand in hand. Additionally, while traditional EEG features might not give rise to the full picture of value's neural correlates, using advanced non-linear methods, provided better prediction accuracies.

P1-B-4 - Brain responses to vehicles predict individual interest and forecast changes in aggregate demand

Dylan Christiano¹, Tara Srirangarajan², Shabnam Hakimi³, Matthew Klenk³, Monica Van³, Charlene Wu², Kalani Murakami³, Brian Knutson²

¹Stanford University & University of Michigan, ²Stanford University, ³Toyota Research Institute

Objective: Demand for vehicles using alternative fuel sources is increasing, leading to the introduction of new vehicle models and types. Little is known, however, about market demand for these new vehicles. We sought to examine whether brain activity could predict individuals' desire to purchase and learn more about vehicles (Erk et al., 2002), as well as forecast changes in demand for vehicles out-of-sample in the US market (Knutson & Genevsky, 2018).

Method: Thirteen subjects who reported being interested in purchasing a vehicle in the next two years participated in a vehicle rating task as their brain activity was monitored using Functional Magnetic Resonance Imaging (fMRI). During each trial of the vehicle rating task, subjects first saw a centrally presented image of the vehicle and its name (2 sec), followed by its fuel source (2 sec), followed by rating scales for interest in learning more and desire to purchase that vehicle (4 sec each).

Individual analyses sought to predict rated desire to know more about each vehicle based on neural data that was averaged and extracted from volumes of interest in the Nucleus Accumbens (NAcc), Medial Prefrontal Cortex (MPFC), and the Anterior Insula (AIns) (Samanez-Larkin and Knutson, 2015). Aggregate analyses sought to forecast change in demand for units sold over 2022 (versus average demand during 2022) using archival data (i.e., goodcarbadcar.com) based on brain activity extracted from the same volumes of interest, but averaged over the sample by vehicle.

Results: Individual choice and whole brain analyses revealed that subjects' initial NAcc response to vehicles (i.e., first 2 sec after a 4 sec lag) predicted they would want to know more about and purchase those vehicles. Activity in other volumes of interest did not significantly predict wanting to know more about or purchase vehicles.

Average NAcc response to vehicles forecast the slope in market demand of units sold ($p < 0.006$), but not average units sold during 2022. Average desire to know and purchase ratings did not forecast the slope or average units sold. In a regression model combining NAcc activity with behavioral ratings, only NAcc activity significantly forecasts slope of units sold.

Conclusions: Brain responses to vehicles, specifically early responses in the NAcc, predict individuals' desire to purchase and learn more about them. Further, average group NAcc activity can forecast changes in demand for cars. These findings suggest that brain activity might add value to conventional measures for forecasting consumer demand for different types of new vehicles.

P1-B-5 - Structural Properties of Association Networks Predict Preference-Based Set Choice

Kianté Fernandez¹, Uma Karmarkar², Ian Krajbich¹

¹University of California, Los Angeles, ²University of California, San Diego

Objective: How do people choose between menus of options, like when choosing between restaurants or shops? Comparing sets of items is complex and may depend on their cohesion; a shop with similar items may be better than one with dissimilar items. A major challenge in addressing such questions is quantifying similarity. Here, we define similarity as the degree to which items are associated with each other in network models, where relations between items (nodes) are represented as connections (edges); items that fit together are closely and densely connected. We use network models to understand how items relate to each other and how those relations affect choices between sets.

Methods: To quantify relations between items, we estimated an association network using a study where subjects rated a variety of food items ($N = 267$; Lee & Holyoak, 2021). We extracted several measures of node importance and connectedness, to see how those features affect choices between sets of those foods. In two new studies ($N = 30$; $N = 75$), subjects chose between two sets of six foods each; the studies varied on whether subjects would receive all the items or only one item from a chosen set. Subjects also provided liking ratings for each food; in the second study they also provided similarity judgments for each set of foods.

Results: Our network analysis yielded 7 communities, which mapped onto categories like fruit, chocolate, and chips. Within communities, foods that would have been predicted to be similar, like raspberries and blackberries, were indeed closely related. We found that network structure played a significant role in choosing between sets (Study 1: $\beta = 0.16$, 95% HDI [0.05, 0.26], Study 2: $\beta = 0.11$, 95% HDI [0.04, 0.18]). Controlling for items' liking ratings, subjects were more likely to choose sets containing items with high node strength and clustering coefficient. Moreover, subjects' similarity judgments correlated with the strength of connectivity between foods in the set ($p = 0.42$, 95% HDI [0.25, 0.55]).

Conclusions: Unlike previous work deriving semantic memory representations from statistical regularities in text (Bhatia & Aka, 2022) or semantic fluency tasks (Zemla & Austerweil, 2018), we show that relational representations can be derived from preference data, and that they align well with subjective measures of similarity. Moreover, we demonstrate that these relations affect people's choices between sets – people prefer sets with more well-connected items. Our work highlights the usefulness of network science in the study of preferences and provides a new tool for assessing how well sets of items fit together.

P1-C-6 - Available Cognitive Resources Shape Subjects' Sophistication in Games

Vered Kurtz-David¹, Adam Brandenburger¹, Paul Glimcher¹

¹New York University

Objective: Finite cognitive capacity limits human reasoning. But how changes in available cognitive resources shape the sophistication of analysis by players in strategic interactions has yet to be unravelled. Here, we investigate how adult humans reason and act in strategic interactions using orthogonal cognitive functions that have previously been associated with distinct cortical modules. Via a simple experimental approach, we show how manipulation of the cognitive demands of a given game affects the elasticities between cognitive modules.

Methods: We took a revealed-preference approach to the analysis of cognitive capacity along two psychological dimensions – social complexity and choice complexity – during game-theoretic interactions. Each game instance was described by the number of social-dependencies, or “iterations of thinking about other players”, and the number of choice options involved. Together, these two properties implied a measure of the total complexity of a given problem. If human players have limited finite overall cognitive capacity to deploy, it follows that in any sufficiently complex game problem, they must trade-off cognitive resources across these two dimensions. At the economic level, the problem creates an imperfect (inelastic) trade-off of cognitive resources that varies as task demands change. To test these ideas, we created a series of variations on the ring game (Kneeland, *Econometrica*, 2015). In this game, each player's payoff depends on their own choice and the choice of the next player in a ring of players. The number of players in each ring determines the social complexity, while the number of options facing each player determines the choice complexity. Further, by varying the exposure time that subjects were given to process information in each strategic interaction, we manipulated the availability of cognitive resources.

Results: We used a classic Cobb-Douglas production function to perform structural estimation of subject-specific cognitive capacities and trade-offs. We found that the elasticities between social and choice demands varied on the subject level, as evidenced in heterogeneity in our model curvature parameter. Average sophistication of subjects was captured by the model constant. As exposure times for the processing of information increased, average sophistication increased logarithmically, in apparent accord with Hick's Law.

Conclusions: Our framework and findings suggest that, even within the same class of games, subjects' sophistication is endogenous with respect to cognitive capacity.

P1-C-7 - The Effects of Arbitrary Social Comparison on Trust and Risky Decision Making

Shulang Yue¹, Iris M Vilares¹

¹University of Minnesota, Twin Cities

Numerous fields have demonstrated that our decisions are guided by social information, and one important piece of social information is social comparison. However, virtually all prior studies investigating the influence of social comparison on trust decisions employed priming and performance based social comparison as comparison direction paradigms. These paradigms are often concerned with personal meanings including one's abilities, skills, or achievements. Nevertheless, we frequently find ourselves in specific social positions for no obvious reasons. Building on the work of Moyal et al. (2020), the current study aims to investigate how trust, reciprocity, and risk related decisions are made subsequently in an unrelated situation after experiencing arbitrarily manipulated social comparison contexts. In addition, the study aims to compare people's behaviors in trust and risk environments under different social positions.

The study employed a 3 (Comparison direction: upward, lateral, downward; between subject) by 2 (Economic game: trust, risk; within subject) mixed design. A power analysis determined that a sample size with at least 252 participants is needed to detect a medium-sized effect ($f = 0.25$) with a 95% power. Participants completed three tasks: a social comparison induction task, a one-shot Trust Game, and a one-shot Risk Game. The social comparison induction task was always completed first, and the order of the Trust and Risk games was counterbalanced.

Our preliminary results (50 participants) showed a significant difference of how decisions were made under trust and risk environments ($d = 0.3$). In general, participants transferred smaller amounts in the Risk Game than in the Trust Game. We found that making downward social comparisons resulted in less subsequent risk-taking behavior, while making upward social comparisons resulted in an increased tendency to take risks. Moreover, we observed an interesting pattern in which participants in the upward comparison condition treated trust and risk environments similarly, while those in the lateral and downward comparison conditions treated them differently. However, no significant difference was observed for the effect of arbitrary-based social comparison on trust and reciprocity decision-making in our preliminary data.

Overall, our preliminary findings suggest that people behaved differently in subsequent trust and risk environments after exposure to randomly manipulated social comparison contexts, and that they perceive risk and trust environments differently under distinct social positions. However, a large sample size is needed to draw firm conclusions.

P1-C-8 - The single- and dual-brain mechanisms underlying the adviser's confidence expression strategy switching during influence management

Shuyi Zha

Effective influence management during advice-giving requires individuals to express confidence in the advice properly and switch timely between the 'competitive' strategy (i.e., give overconfident advice) and the 'defensive' strategy (i.e., give underconfident advice). However, how advisers switch between these two strategies, and whether and why there exist individual differences during this process remain elusive. We used an advice-giving game that manipulated incentive contexts (Incentivized/Non-Incentivized) to induce the adviser's confidence expression strategy switching and measured the brain activities of adviser and advisee concurrently using functional near-infrared spectroscopy (fNIRS). Sixty-eight Chinese students were recruited (34 pairs) in same-gender, unfamiliar pairs and were then randomly assigned to play the roles of adviser and advisee in the present study. Behaviorally, our results showed that, in general, the advisers applied 'defensive' strategies when they were incentivized (average $\beta_{\text{selection}} = 0.10 > 0$, $t = 2.07$, $p = 0.047$) and 'competitive' strategies when they were not incentivized (average $\beta_{\text{selection}} = -0.09 < 0$, $t = -1.77$, $p = 0.087$), and demonstrated significant differences in $\beta_{\text{selection}}$ across incentive contexts ($t = 2.04$, $p = 0.045$). Moreover, we observed the individual differences in strategy switching in different incentive contexts. The I-D_N-C group was more effective in influence management than the I-C_N-D group, as reflected by the frequencies of advice-taking in the advisees ($r = 0.41$, $p = 0.022$). Neurally, compared with the I-C_N-D group,

the I-D_N-C group showed higher activation in the dorsolateral prefrontal cortex (DLPFC) that supported strategy switching (CH4; $t = 5.91$, $p = 0.011$, FDR corrected), as well as increased interpersonal neural synchronization (INS) in the temporoparietal junction (TPJ) that supported influence management (CH13, $t = 2.21$, $p = 0.048$, FDR corrected). This two-in-one process, i.e., confidence expression strategy switching and the corresponding influence management, was linked and modulated by the strength of DLPFC-TPJ functional connectivity in the adviser ($\beta a = 0.45$, $SE = 0.03$, $t = 1.76$, $p = 0.037$; $\beta b = 0.46$, $SE = 0.03$, $t = 2.55$, $p = 0.021$; $\beta c = 0.47$, $SE = 0.03$, $t = 2.84$, $p = 0.016$). We further developed a neurocognitive model that contributed to understanding the adviser's strategy switching during influence management. This work provides a theoretical base for future investigations on strategic advice in various social contexts.

P1-C-9 - How do adolescents sustain cooperation with peers versus highly adaptive computer partners?

Wenda Liu¹, Isabella Mark¹, Christoph W. Korn², Gabriela Rosenblau¹

¹George Washington University, ²Heidelberg University

Cooperation is a crucial prosocial skill that develops during adolescence and relies on trust in our interaction partner. While traditionally studied in dyads or small groups of conspecifics, the emergence of adaptive artificial intelligence (AI) systems raises the question of whether humans differ in their cooperation decisions and strategies when interacting with human versus AI partners.

The study examined how adolescents cooperate in economic exchange games with real peers (real interaction) versus adaptive computer algorithms. To test the effects of prior beliefs, we told participants on some of the games that they kept playing with their peer even though they played with an adaptive computer (social condition). On other games, we told them that they played with a computer (non-social condition). Our preregistered hypothesis was that adolescents would be more generous in the real interaction and social condition as supposed to the non-social condition.

Adolescents ($N = 35$, mean age = 11y ($\pm 2.6y$, 19 males) played a child friendly version of the multi-round trust game (26 came in pairs, matched for age, sex and IQ). Participants played five consecutive games. Pairs first played the trust game with each other (real condition). To this end, partners were randomly assigned to play the investor or trustee. Thereafter, each participant played four games in the investor role with the computer algorithm. Before two of the four games, participants were told that they were playing with their peer (social condition). In the remaining games participants were told that they were playing with computer partners (non-social condition). The order of conditions were counterbalanced.

Adolescents were more generous during the real interaction and the social compared with the non-social condition ($\hat{\rho}^2 = 0.11$, $SE = 0.03$, $p < 0.001$, Cohen's $d = 0.38$). This was true even though the trustee reciprocated more often in the social and non-social conditions as supposed to the real peer interactions. Despite this, adolescents reinitiated cooperation more often after the trustee defected in the real interaction as supposed to both social and nonsocial ($\hat{\rho}^2 = 0.23$, $SE = 0.09$, $p < 0.001$, Cohen's $d = 0.67$). To test participants' strategies during games, we fit a reinforcement learning model and a fixed economic strategy (tit-for-tat) as well as a combination between these two. We found that adolescents were more cooperative when playing and believing that they played with a peer than with an adaptive computer. Participants followed an economic strategy, tit-for-tat, to sustain cooperation across conditions.

P1-D-11 - Predicting Entrepreneurial Status: Insights from Gray Matter Volume and Brain Responses to Uncertainty

Tom Lahti¹, Kaisa Hytönen², Joakim Wincent¹, Iiro Jääskeläinen³, Martin Schürmann⁴, Marja-Liisa Halko⁵

¹Hanken School of Economics, ²Laurea University of Applied Sciences, ³Aalto University School of Science, ⁴University of Nottingham,

⁵University of Helsinki

Objective: Entrepreneurs are known for their higher tolerance of risk and uncertainty, which can play a crucial role in their decision to become entrepreneurs. To explore if entrepreneurs respond differently to risk and increasing ambiguity compared to others, we conducted a comparative analysis of behavioral responses, brain activation patterns, and gray matter volume between entrepreneurs and employees without entrepreneurial backgrounds. Additionally, we employed a machine learning approach to predict participants' entrepreneurial status based on our fMRI and GMV data.

Methods: We analyzed functional and structural MR images along with behavioral data from 20 entrepreneurs and 20 employees. We used the Becker-DeGroot-Marschak (1964) mechanism to evaluate how entrepreneurs and the control group valued risky and ambiguous options (lotteries). Participants assigned values through willingness-to-sell prices. Utilizing fMRI data, we identified brain regions showing different responses to risk and ambiguity between entrepreneurs and employees, while GMV data helped estimate individual risk attitudes. For predicting participants' entrepreneurial status, we employed a nested cross-validation procedure.

Results: Entrepreneurs value lotteries significantly higher than the control group and are less affected by ambiguity compared to employees. Individual characteristics like risk attitude, affect intensity, optimism, or confidence do not directly explain the difference in valuations and attitudes toward uncertainty. In fMRI data analysis, we found significant differences between the groups in the encoding of value during risky trials in the bilateral cerebellum and during ambiguous trials in the right cuneus and ACC. The GMV data analysis indicates a highly significant negative association between DMPFC and participants' risk attitudes. Risk-taking propensity, especially among entrepreneurs, shows a positive relationship with GMV in the right and left anterior insula. For predicting participants' entrepreneurial status, the model integrating valuation decisions with fMRI ROI data during risky trials achieves the highest predictive accuracy (85%). Particularly, participants' brain response in the left anterior insula and left cerebellum during risky trials demonstrates a positive association with entrepreneurial status.

Conclusions: Our results suggest that entrepreneurs' heightened risk tolerance might be rooted in their brain activation patterns. The predictive power of our models, achieving 85% accuracy, demonstrates the potential utility of fMRI and GMV data in understanding and identifying entrepreneurial behavior.

P1-D-12 - Acute stress induces risk-seeking via more optimistic beliefs

Gilles De Hollander¹, Maike Renkert¹, Saurabh Bedi¹, Gökhan Aydoğan¹, Christian Ruff¹

¹University of Zürich

Acute stress is an inevitable aspect of life, with long-lasting consequences for physical and financial well-being (DeLongis et al 1988). Previous research documents a link between acute stress and altered risk preferences (Buckert et al. 2014), with suggestions that stress may contribute to the perpetuation of poverty (Haushofer & Fehr 2014). However, despite ample research, little is known about the neurocognitive processes that translate stress into altered risky decision-making.

Here, we shed light on these processes in an experimental study of risky choices under laboratory-induced stress. We employed a perceptual account of risky choice assuming that decision-makers make financial decisions based on noisy and biased perceptual representations of payoffs (Khaw et al., 2020; Garcia et al. 2023). This allowed us to decompose stress-induced shifts in risk preferences to latent Bayesian perceptual processes, specifically to either noisier sensory representation or altered beliefs.

Participants (n=50) performed a risky gamble task in a first baseline fMRI session, before being randomly assigned to either a stress or control group. To induce stress in the second fMRI session, we utilized an adapted version of the well-established MIST task based on social evaluative threat. Cortisol measures across 6 timepoints during the session confirmed that the stress induction was successful, as cortisol levels were significantly higher in the stress group compared to the control group ($p=2.5 \times 10^{-6}$).

Stress led to a systematic shift in risk preferences, with more risk-seeking behavior under stress ($p=0.03$). Systematic model comparisons showed that both behavior in the first session and the context-dependent shift in the stress session were best explained by our new perceptual model, as compared to classic utility models of risky choice (e.g. EUT). Inspection of model parameters revealed that the shift in risk preference under stress was induced by more optimistic (and mostly more realistic) prior beliefs about the magnitude of the risky payoffs ($p<0.01$). By contrast, the noisiness of the inferred payoff representations was unaltered.

In conclusion, our Bayesian perceptual approach allowed us to provide a mechanistic perspective on the cognitive processes involved in stress effects on risky choice. Our results indicate that stress does not lead to noisier processing of information, but to shifts in prior beliefs. In upcoming work, we will link model parameters to the neural data, to neurally validate the model-inferred processes and to elucidate how the corresponding neural mechanisms are affected by stress.

P1-D-13 - Utility function changes with the gamble's winning probability

Yilong Lu¹, Yang-Fan Lu¹, Hang Zhang¹

¹Peking University

Objective: Recent studies suggest that people's utility function may not be as stable as traditionally believed but may vary with attention or decision context in a dynamic way. However, it is still widely believed that value and probability are two independent dimensions, with the utility function of a value not depending on the probability associated with it. Here we conducted a series of online experiments of decision under risk to further advance the dynamic view, investigating whether and how the utility function may be influenced by the gamble's winning probability.

Methods: On each trial, participants chose between a two-outcome gamble (p, x_1) and a sure reward c . The gambles in each experiment were generated from a factorial combination of various values (x_1) and winning probabilities (p). We developed both nonparametric and model-based methods to measure participants' utility function for each different probability. These experimental designs and data analysis methods allowed us to obtain measures without relying on particular assumptions of probability weighting functions as well as to minimize known artefacts in utility measurement. To verify the robustness of our findings, we varied the way of x_1 - c pairing, the viewing (decision) time, and the range of p across four experiments (total $n=175$).

Results: In all four experiments, the nonparametric and model-based analyses led to the converging finding that participants' utility function was less concave (i.e., less risk-averse) when the winning probability of the gamble was higher. According to Bayesian parameter estimation in the model-based analysis, the measured value sensitivity term α (exponent of the power utility function) increased with p in all four experiments (95% HDI of group-level increment larger than 0). The difference was remarkable: for $p=0.3$ vs. $p=0.9$, the mean α was 0.47 vs. 0.87 (Exp 1), 0.47 vs. 0.69 (Exp 2), and 0.65 vs. 1.25 (Exp 3); from $p=0.01$ to $p=0.9$, the mean α increased from 0.29 to 2.34 (Exp 4). The findings that α increased with viewing time (Exp 4 vs. Exp 2 & 3) as well as with the winning probability is consistent with the resource-rational (rational-inattention) hypothesis of value coding, implying that limited cognitive resources are more likely to be allocated to process more probable outcomes.

Conclusion: As the gamble's winning probability increases, the utility function curve steepens. Our finding further challenges traditional static views on utility functions and reveals a cross-dimensional interplay in the processing of value and probability that indicates resource rationality.

P1-D-14 - Survey and Incentive-Compatible Approaches: Risk Measures, but Not Discounting Measures are Correlated.

Zih-Yun Yan¹, Paul Glimcher¹

¹New York University

Preferences have been studied by both psychologists and economists, although using disparate methods. While both surveys and incentive-compatible tasks aim to capture the same constructs, recent studies have shown only a weak correlation between these methods. Two explanations for the low correlation have been proposed – that the two methods measure different mental processes or the low test-retest reliability of behavioral tasks obscures an underlying correlation. To determine which is correct for risk, ambiguity, and delay discounting preferences, we developed a new experimental protocol to enhance test-retest reliability. Using this approach, we reexamined the correlation between survey and task measures of these important preferences. Our major improvements were the use of multiple incentive-compatible measurements in a longitudinal experiment with a well-stratified sample. We used a standard lottery task to study risk and ambiguity preferences and a standard delay discounting task. We compared these to the standard Domain-Specific Risk-Taking Scale (DOSPERT), the standard Intolerance of Uncertainty Scale (IUS), and the standard Urgency-Premeditation-Perseverance-Sensation Seeking (UPPS) impulsivity scale. We conclude that the DOSPERT does measure the same underlying mechanism as behavioral risk and ambiguity tasks, at least in young adults. However, we find no evidence that the mechanisms accessed by the standard Intolerance of Uncertainty (IUS) are connected to those accessed by incentive-compatible lottery choice tasks. Examining time preferences, we found that the standard UPPS scale for measuring impulsivity was completely uncorrelated with intertemporal impulsivity as measured with a standard discounting task. While surprising, this suggests that financial impulsivity, measured with a discounting task, and behavioral impulsivity, measured with the UPPS, reflect fundamentally different constructs.

P1-E-15 - Retrospective bias in goal pursuit

Sneha Aenugu¹, John O'Doherty¹

¹California Institute of Technology

Objective: The choice to stay committed to a temporally-extended goal or switch away from it entails weighing retrospective value (how much has been accomplished so far) against prospective value (how much further till the finish line). It appears that people often stick with an option longer than is viable, a phenomenon documented as retrospective cost or sunk cost in decision-making (Arkes & Blumer, 1985). On the other hand, contemplating the prospective cost – future costs and benefits of further investment – is considered a rational approach. In a novel task where an option needs to be persistently executed till a set target to earn rewards, we demonstrate an undue bias in favor of retrospective value in human behavior resulting in sub-optimal performance. We also formulate the computational mechanisms to account for this bias.

Methods: Subjects (n=70) participated in a suit collection paradigm where they collected suits of tokens of three different types (7 tokens of the same kind constitutes a suit) to earn rewards. Subjects played 540 rounds of the game which were grouped into 18 blocks with shifting token contingencies. Each block is one of three types: 80-20, 70-30, 60-40. For instance, in the 80-20 condition one token is received with 80% chance and the other two are available with 20%. The most abundant token switches across adjacent blocks.

Results: Embedded in the framework of reinforcement learning, we introduce two models: prospective and retrospective. Prospective model explicitly encodes token contingencies and uses it to estimate the value of staying committed to a token at every instant, while retrospective model performs the task reactively and is biased in favor of tokens that are partway through the target. We demonstrate that human behavior shows influences of both prospective and retrospective computations and the degree of prospective valuation in individuals is robustly correlated with task performance in each condition: $R=0.87$ (80-20), 0.76 (70-30), 0.67 (60-40). We fit a model that employs a weighted arbitration between prospective and retrospective models and show that the weight parameter on prospective value fit on subject-level robustly correlated with task performance ($R = 0.68$, $p < 0.0001$).

Conclusions: Our study introduces a novel paradigm to demonstrate the bias in temporally extended goal-valuation in humans and proposes computational hypotheses to account for this bias. Therefore, our study opens up a new avenue to study the manifestations and consequences of cognitive biases in everyday goal pursuit.

P1-E-16 - Pain now! Exploring the relationship between pain, temporal uncertainty, and inter-temporal choices

Mathieu Roy¹, Taryn Berman¹

¹McGill University

Making decisions about future pain is essential, as poor decisions can lead to unnecessary pain. A type of bias extensively studied with rewards, but rarely with pain, is delay discounting (i.e., the preference for smaller rewards sooner over larger future rewards). Studies using hypothetical pain or short delays have found a similar effect with pain, where people accept more pain in the future over less pain now, as it decreases in aversiveness across time (Harris, 2012). However, others suggest that people prefer pain now to get it over with (i.e., dread; Story et al., 2013). Across two studies, we sought to investigate how inter-temporal choices for real pain are made, the influence of temporal uncertainty, and what causes these inter-temporal differences. In the first study, 60 participants performed an inter-temporal choice task wherein they selected between two choices which differed in pain intensity and delay (i.e., 15s, 30s, 1-hour, and 1-month). Importantly, participants were always choosing between pain now and at delay. In the second study, 80 participants performed the same task with a temporal uncertainty modification. In this between-groups design, one group – the certainty group – knew exactly when the pain would occur, while the other – the uncertainty group – did not. In the first study, we found that subjects accepted more pain in order to experience it sooner and had an overall preference for pain now ($d=1.41$). Using multilevel modeling, we found that delay and pain intensity for each choice offered interact to predict pain now decisions ($R^2=.78$). Interestingly, when computational modeling was used, the data fit a typical delay discounting curve where delayed offers were viewed as worse. The second study replicates the first, where both groups displayed a preference for pain now ($dcertainty=1.15$, $duncertainty=1.64$), and this preference was greater for the uncertainty group. Multilevel modeling revealed that group, delay, pain levels for each choice, and intolerance of uncertainty all predict pain now decisions ($R^2=.82$). Similar computational modeling results were found for both groups, such that delayed offers were viewed as worse options, regardless of the pain levels. Overall, people would rather accept more pain now than wait for less, suggesting that delay is aversive. This effect may be partly due to uncertainty about when the pain will occur and one's tolerance to uncertainty. This study is the first of its kind to examine inter-temporal choices with real pain and has important implications for interventions aimed to reduce detrimental biases that lead to unnecessary pain.

P1-E-17 - "Will" it be Seen? Using Eye Tracking to Re-examine the Future Tense Effect in Intertemporal Choices

Josie Chen¹, Wei Chen¹, Jiang-Shiang Hu¹

¹National Taiwan University

The Linguistic-Savings Hypothesis (LSH) proposes that using future tense references can decrease future-oriented behaviors. We conducted an eye-tracking experiment to investigate the cognitive processes and causal effect of the hypothesis, which has previously only been studied for its behavioral outcomes. To test our hypothesis, we used modified multiple price lists. Each problem presented a choice between a smaller payment sooner or a larger payment later. The linguistic treatment was embedded in the description of the larger payment. For the treatment group, the description was "After Y months will receive Z tokens". For the control group, the description was "After Y months, receive Z tokens" without the future tense "will". We recruited 73 subjects and randomly assigned them to either the treatment group (36 subjects with future tense) or the control group (37 subjects without future tense). Our results showed that the treatment group, who saw the "will receive" phrasing, chose 12.5% fewer later-larger payments compared to the control group. This result was consistent with the LSH but the effect size was not significant in the logistic regression model ($p\text{-value} = 0.366$). Additionally, we found that the treatment group took 8.8% longer to make their decisions compared to the control group, which was significant in the linear regression model ($p\text{-value} < 0.001$). We used a modified Drift-Diffusion Model (DDM) to explore the cognitive processes involved. Our findings suggest that the future tense has three effects on decision-making. Firstly, it makes the later-larger reward appear more distant, which increases the preference for the sooner-smaller payment. This is reflected in the drift rate parameter "wt" ($p\text{-value} = 0.001$). Secondly, the future tense raises the decisional threshold "a" ($p\text{-value} = 0.038$), indicating that people approach intertemporal choices with greater caution. To study the impact of the future tense, we used group parameters from the DDM to simulate hypothetical subjects completing 112 decision problems. The treatment effect of the future tense becomes

significant (p -value < 0.05) when more than 600 subjects complete the task. This suggests that the null result in the logistic regression was due to a small sample size. In other words, the future tense affects both decision-making and outcomes. To our knowledge, this is the first paper to uncover the cognitive/computational process of LSH using a neuroscience approach. Our paper reconciles the conflicting results in the LSH literature and sheds light on the importance of modeling the dynamics of the choice process.

P1-E-18 - Understanding the mechanism of perceived control: the role of self-efficacy and response-efficacy

Yun-Yen Yang¹, Mauricio Delgado¹

¹*Rutgers University - Newark*

An individual's perception of how much control they have in the environment can exert a significant influence on everyday decisions. Indeed, participants often prefer situations where they can exert their choices (e.g., picking an airline seat) over not being in control (e.g., letting the computer assign them a seat) even if they have to incur a cost to maintain control - a process thought to be mediated by the ventromedial prefrontal cortex (vmPFC; Wang et al., 2019). Critically, perceived control can be distinguished into two components (Bandura, 1977): self-efficacy, referencing the belief of an individual in successfully executing a behavior, and response-efficacy, regarding the belief that the behavior itself will lead to an expected outcome. This suggests two potential pathways through which perceived control can influence decisions; however, prior research has not fully considered this distinction. In the study, we investigate the neural representations that differentially code for self-efficacy and response-efficacy, and how such signals may be integrated to inform decision-making. Participants completed a novel task where they had to make a perceptual decision about whether they could hit a moving target on the screen (self-efficacy), or whether the presented card that highlighted the probability of that trial being worth a reward (response-efficacy) would win. We scanned participants ($n=39$) while they were presented with either one or two pieces of information simultaneously and asked them to rate their confidence in a) performing the behavior (self-efficacy), b) the potential of a rewarding outcome (response-efficacy), and c) successfully playing the game (integration). Computational modeling was used to measure the subjective weight of self-efficacy and response-efficacy while making decisions and used to investigate the neural mechanism of integration. Participants incorporated both self-efficacy and response-efficacy information while making choices, instead of purely using expected value. At the neural level, activity in the vmPFC was positively correlated with self-efficacy, whereas activity in the striatum positively correlated with response-efficacy. Additionally, the subjective value of playing the game was found to be represented in the dorsal anterior cingulate cortex (dACC), supporting a possible information integration role, consistent with prior observations of dACC in the allocation of control (Shenhav et al., 2016). Taken together, these findings further our understanding of how perceived control impacts decision-making, highlighting different pathways that influence our perception of control.

P1-F-19 - Effects of social vs. non-social stimuli on threat learning in Pavlovian-to-instrumental transfer

Nidhi Desai¹, Sophie Paolizzi¹, Michael Hallquist¹

¹*University of North Carolina, Chapel Hill*

People take information learned from others into account when choosing optimal actions (Charpentier et al., 2020). However, overutilizing social information can undermine effective decision-making (FeldmanHall & Shenhav, 2019; Lamba et al., 2020). Recent work suggests socially transmitted information biases instrumental choice via the Pavlovian system (Baczkowski et al., 2023; Lindström et al., 2019). While such social Pavlovian-to-instrumental transfer (PIT) might facilitate effective information transmission, detecting and overriding social Pavlovian biases when social interactants are not providing accurate information is critical for optimal instrumental action selection. We tested whether adaptability in instrumental decisions was modulated by social and non-social Pavlovian stimuli using a new task designed to study avoidance of an aversive outcome via PIT. 51 individuals, either endorsing psychopathology relevant to social dysfunction (e.g., borderline personality disorder or social anxiety) or low in clinical symptoms, underwent Pavlovian conditioning on social and non-social cues (6 trials each) predicting either a loud thunder sound or no thunder outcome. Then, participants made instrumental choices between a pair of these stimuli (60 trials) to avoid the aversive outcome. The instrumental relevance of Pavlovian associations was manipulated using a contingency reversal (1) partway through the decision phase (25th-35th trials) and (2) in initial transfer phase for half the blocks, partway switching back. We used multi-level logistic regression to evaluate how accurately participants learned to use or discard Pavlovian information provided by both cue types (social/non-social). Learners were less accurate in selecting stimuli to avoid thunder when the Pavlovian association was misleading about the instrumental contingency. Accuracy also suffered more when reversals occurred following accurate compared to inaccurate transfers, suggesting performance suffered when Pavlovian biases were incongruent with an instrumental contingency. People were less accurate when stimuli predicting aversive outcome were social compared to non-social, yet they experienced less severe drops in accuracy when contingencies changed, and social cues were no longer predictive of aversive outcomes. Altogether, our findings, consistent with previous work on the associability of social stimuli, suggest that Pavlovian social stimuli strongly influence instrumental decisions via PIT and their adaptability to changed environment. Funding: National Institute of Mental Health R01MH119399 (PI: Michael Hallquist)

P1-F-20 - The Influence of Vasopressin Deficiency on Social Decision-Making in CDI Patients

Mikhail Votinov¹

¹*Research Center Jülich*

The peptide hormone arginine vasopressin (AVP) is critical for maintaining homeostasis and also plays a pivotal role in complex social decision-making and behavior across various species. Damage to AVP-secreting neurons can cause central diabetes insipidus (CDI, prevalence of 1-9/100,000). We hypothesize that CDI patients exhibit comparable social deficits reported in animal models with AVP deficiency. This study explores impaired social decision-making in CDI patients compared to healthy controls through analysis of their prosocial and approach/avoidance behavior. Blood samples were taken to measure copeptin, a surrogate of vasopressin. The study involved 12 CDI patients and 12 healthy controls (HCs), with equal gender distribution and age range. Participants engaged in three games to assess their social behavior and decision-making: the social-discounting task (SDT) to evaluate prosocial and altruistic tendencies, the ultimatum game (UG) to measure their perception of fairness and the Stop Distance Task to determine their approach towards a person/animal in three emotional states (friendly, neutral and aggressive). The hormone analysis revealed a substantial decrease in copeptin levels in the CDI group compared to the HC group, which confirms the presence of AVP deficiency. The findings of the distance task demonstrated that the approach distance (AD) of participants was influenced by the emotional expression of the confronting character with a decreasing distance from aggressive to neutral to friendly across all subjects. CDI patients compared to HC chose a significantly larger AD between themselves and human subjects across all emotional states, except for friendly female and

neutral male. Notably, there were no significant differences between the CDI and HC in the dog condition across all emotional states. Furthermore, the results of the SDT showed increased generosity in CDI patients, but not statistically significant. The UG also did not yield significant results.

In comparison to HC, patients with CDI displayed reduced approach behavior towards human characters, regardless of their emotional expression. However, this effect was observed specifically in condition with humans but not with animal characters, which may suggest that CDI patients may experience higher levels of social anxiety. Although there was a trend towards increased generosity in CDI patients, the results were not statistically significant, possibly due to the limited sample size. Further research is needed to extend these findings.

P1-F-21 - Effort Allocation in Competition and Cooperation

Nicholas Angelides¹, Michael Platt¹

¹University of Pennsylvania

Models of social decision-making historically rely on monetary costs to the self in computations of utility, or subjective value. Both cooperative and competitive behavior, measured using social economic games, are standardly modeled on preferences regarding monetary outcomes. However, many social behaviors, in which an agent might incur costs to themselves to help or hinder another, regularly entail a primary biological cost: energy. Social acts in both competition and cooperation are an evolutionarily adaptive and crucial component of human social coexistence, and in real world scenarios often require the allocation of effort. The way in which such metabolic costs affect the valuation of a decision with shared (cooperative) versus self-rewarding and other-punishing (competitive) is understudied. This study adapts a well-studied effort-based decision-making tasks and employs computational modeling to examine how competitive versus cooperative environments modulate willingness to engage in an effortful task. By parametrically modulating reward and effort within two social environment conditions per participant, we will examine both population-wise and within-subjects variation in choice behavior when faced with working with versus against a conspecific. To formally characterize such variation, we will employ computational models of effort-based decision-making to capture and predict effort discounting behavior as a function of either competition or cooperation. Using physical effort as a cost manipulation enables us not only to calibrate precise cost levels per participant, but crucially allows us insight into a more primary, biological cost to self: metabolic energy. Results from this study will shed light on how the influence of social environment modulates some of the same energetically expensive but evolutionarily adaptive behaviors preserved in modern social life.

P1-F-22 - Linking Brain-Wide Activity Patterns during Neuroeconomic Decision Making to Aggression

Romain Durand-De Cuttoli¹, Brian Sweis¹, Antonio Aubry¹, Long Li¹, Eric Nestler¹, Scott Russo¹

¹Icahn School of Medicine

Aggression is an evolutionarily conserved response to threat spanning a range of behaviors, including those that may be adaptive versus those that may be pathological and dangerous to others. Circuits recruited during aggressive behaviors, in addition to circuits known to subserve social interactions themselves, are critical nodes of the brain's reward system. However, it is unclear how pathological aggression may engage key reward circuits in the brain, contributing to maladaptive reinforcement of aggression. Because the brain has evolved to use multiple decision-making systems, simple tests of reward value may be unable to access computational subtleties that may be typified by aggressive individuals.

We characterized decision-making profiles of 40 outbred Swiss Webster mice screened for aggression and then tested on the neuroeconomic task, Restaurant Row. Mice had limited time each day to forage for their sole source of food investing in rewards of varying costs (delays from 1-30s signaled by tone pitch) and value (unique flavors tied to four spatially cued locations). On the final day of testing, mice engaged the task before being prepped for whole brain iDISCO+ tissue clearing and staining in 275 distinct brain regions for c-Fos expression, an activity-dependent immediate early gene.

Overall, we found that the majority of brain regions revealed decreased levels of c-Fos expression in highly aggressive animals versus non-aggressive animals. Using an unbiased, open-ended analysis approach, top region hits revealed strong correlations between aggression score and c-Fos expression – regions that lie in the medial wall of the prefrontal cortex (mPFC) and are known to be engaged by the Restaurant Row task. Using this open-ended approach, we also found that several regions across the limbic system covaried with numerous key metrics from the Restaurant Row task, including individual differences in the economic profiles of different types of choices. We found that thresholds of willingness to wait as a function of cost and subjective value varied across levels of aggression but also depending on choice processes engaged. These behaviors differentially scaled with c-Fos expression along the dorsoventral axis of the mPFC.

Our data reveal how brain-wide studies of aggression may reveal changes in circuits affecting only certain types of decisions being processed. These findings set the stage for future experiments manipulating circuit-specific computations, including within functional subregions of the mPFC, in order to augment dissociable valuation algorithms that may be differentially engaged in the aggressive brain.

P1-F-23 - Friends diverge while strangers align: Using fMRI hyperscanning to investigate social interaction in real-time conversations

Sebastian Speer¹, Emily Falk², Laëtitia Mwilambwe-Tshilobo³, Diana Tamir¹, Shannon Burns⁴, Lily Tsoi⁵

¹Princeton University, ²University of Pennsylvania, ³Princeton University & University of Pennsylvania, ⁴Pomona College, ⁵Caldwell University

Successful social interactions are vital, but the key ingredients for a successful interaction are not yet clear. One crucial ingredient for success might be mentalizing - the tendency to consider what others think and feel. People use three dimensions, namely social impact, rationality and valence, termed the 3D mind model, to represent mental states (MS). Previous work on social interactions suggests that linguistic, behavioral, physiological, and neural alignment are associated with positive social outcomes. We therefore tested whether people align their MS over the course of a live conversation, and whether alignment differs between friends and strangers.

We used fMRI hyper-scanning: 60 dyads engaged in a real-time conversation with discrete prompts and demarcated turns. Half of the recruited dyads were friends; The other half were strangers. This allowed us to explore how an already existing social connection influences the MS alignment of dyads during their conversations.

To measure MS alignment, we developed predictive models that can decode MS from whole-brain neural patterns. We used four previous fMRI data sets that used MS judgment tasks designed to evoke neural patterns varying across the three MS dimensions.

These models successfully captured mentalizing during real-time conversation, allowing us to decode from each person's brain their "location" on each dimension in MS space at a given moment of time. We computed the distance between the two speakers in 3D MS space, at each timepoint across the whole conversation, where a smaller distance represents a higher MS alignment within the dyad. We found that strangers' MS aligned over time, whereas they diverge for friends. Friends start off with higher MS alignment than strangers and then drift apart in MS space until their distance is significantly larger than the distance for strangers. Employing NLP we also decoded participants' MS locations based on the words they chose, and replicated the neural findings in linguistic 3d space. Using topic modeling, we found that, in line with divergence in MS, friends explore a larger content space: they generate more topics that are also more distant from each other in content space and switch more frequently between them. This suggests that friends may explore a more extensive content space and exhibit more surprising turns, whereas strangers aim to find common ground in their conversations. Lastly, strangers who diverge in topic as well as MS space also rated the conversation as more enjoyable. Conclusively, social context shapes the extent to which we aim to find common ground or explore new frontiers in a conversation.

P1-G-24 - Delay discounting and risk tolerance in obsessive-compulsive disorder: results from the Global OCD study

Karolina Lempert¹, Benjamin Huber², Marcelo Batistuzzo³, Kartik Sheshachala⁴, Dianne Hezel², Niels De Jooze⁵, Clara Marincowitz⁶, Christine Lochner⁶, Dan Stein⁷, Janardhanan Narayanaswamy⁴, Odile Van Den Heuvel⁵, Melanie Wall², H. Blair Simpson²

¹Adelphi University, ²Columbia University, ³Universidade de São Paulo, ⁴National Institute of Mental Health and Neurosciences, ⁵Vrije Universiteit Amsterdam, ⁶Stellenbosch University, ⁷University of Cape Town

Individuals with obsessive-compulsive disorder (OCD) have shown impairments in reward processing and difficulty with decision making, but previous research examining economic decision making in OCD has yielded inconsistent results. Here we examined, in a large, global sample, whether adults with OCD differ from healthy controls (HC) on two measures of decision making about potential rewards: delay discounting and risk tolerance. We also examined whether clinical measures within the OCD group were associated with these preferences. We analyzed data from the Global OCD study (Simpson et al., 2020; Batistuzzo et al., 2022), a multi-center study of adult medication-free patients with OCD. Standardized versions of delay discounting and risk tolerance tasks were used to collect data in five countries: Brazil, India, the Netherlands, South Africa, and the USA. In the delay discounting task, participants made a series of choices between smaller monetary amounts they could receive immediately, and larger amounts available after a delay. In the risk tolerance task, participants chose between smaller certain monetary amounts and larger amounts available with a 50% chance. Individual discount rates were quantified by fitting a hyperbolic model to data and deriving a discount rate parameter, k . Risk tolerance was quantified by fitting a power-law model to data and deriving a risk tolerance parameter, $\hat{\pi}$. After controlling for age, education level, socioeconomic status, intelligence, and site, individuals with OCD ($n=268$) did not differ from HC ($n=256$) in delay discounting ($\hat{\pi}^2=-0.005$; $p=0.970$) or risk tolerance ($\hat{\pi}^2=0.006$; $p=0.898$). However, patients with OCD who reported more anxiety on the Hamilton anxiety scale showed higher discount rates ($\hat{\pi}^2=0.042$; $p=0.004$), suggesting increased impatience among more anxious patients. Hamilton depression scale scores were also associated with discounting in the OCD group ($\hat{\pi}^2=0.047$; $p=0.010$). No clinical measures were associated with risk tolerance in OCD. Previous studies on economic decision making in OCD have been inconclusive, due to small sample sizes and inconsistency in the inclusion of covariates, such as socioeconomic status. Here, by analyzing a large global dataset of well-characterized patients and controls, we can conclude that OCD patients do not differ meaningfully from HC participants in either of these decision tendencies. However, the finding that symptoms of anxiety and depression within the OCD group were associated with increased delay discounting suggests that anxiety may be associated with more myopic decision making, perhaps trans-diagnostically.

P1-G-25 - The "motive cocktail" in third-party punishment and helping

Xiaoyan Wu¹, Xiangjuan Ren², Chao Liu¹, Hang Zhang²

¹Beijing Normal University, ²Peking University

Objective: Prosocial motives such as social equality and efficiency have been recognized as essential for altruistic behaviors. Yet, the prediction of diverse altruistic behaviors in different situations and individuals remains challenging if one only focuses on 1 or 2 motives, as most previous studies did. Here we aimed to disentangle the potentially numerous and interdependent motives in altruistic third-party punishment and helping behaviors.

Methods: Participants played an unaffected third-party to a dictator game, observing the dictator allocating more of the money to themselves than to the receiver. On each trial, participants were asked to choose whether they would spend a specific amount of money to punish the dictator (or help the receiver) or do nothing. By varying whether the scenario was punishment or help, the level of inequality between the dictator and the receiver, the third-party's intervention cost, and the impact ratio of the payoff change on others to the cost of the third-party, we created 100 different conditions, with each participant completing each condition for 3 times. We initially tested 157 participants in a laboratory study and subsequently replicated the findings with 1258 participants in a pre-registered online study. Such rich datasets allowed us to disentangle the effects of multiple motives as well as their interactions in individual participants' decisions.

Results: Besides the single variable effects that are consistent with previous studies, we found an interaction of inequality \times cost \times ratio (linear mixed-effects model, $b = -0.21$, $p < 0.001$), which indicates that participants might have a more complex utility calculation in their decisions than simply linear combining the variables. Our computational modeling analysis showed that participants' behavioral patterns could not be fully explained by any previous models and were instead best fit by a model we constructed with 7 prosocial motive terms, including 2 "compound motives" that represent interactions between more elementary motives. The interactions suggest that when the intervention cost is higher, the effect of inequality diminishes, a phenomenon we refer to as "inequality inattention". We further identified 3 types of individuals: justice warriors who aim to reduce social inequality, generous helpers who focus on helping the victim and rational moralists who ignored inequality at rising cost.

Conclusions: Our findings expand the rational framework of altruistic behaviors by demonstrating that various motives collectively drive the altruistic behaviors of each individual. These motives do not operate independently but interact with each other.

P1-G-26 - Examining the associations between socioeconomic status and temporal discounting across monetary, social, and health rewards

Christian Benitez¹, Michelle Mahecha Perez¹, Joe Kable¹

¹University of Pennsylvania

Objective: Several studies have shown people from lower SES backgrounds discount the future to a greater extent, and multiple theoretical models have been proposed to account for this phenomenon. One account focuses on the increased financial constraints that lower SES individuals face, including limited savings and access to borrowing options. Consequently, this leads to a prioritization of short-term goals to meet daily financial needs. Another perspective emphasizes lower SES individuals' exposure to early life stressors, leading them to prefer options that minimize the cost of uncertain and unstable future environments. Discriminating between these models is challenging because prior research has solely focused on monetary rewards, while the accounts differ primarily in the extent to which they would predict higher discounting across a broad domain of potential rewards. To address this gap in the literature, this pre-registered study examines the relationship between SES and intertemporal discounting across monetary, social, and health rewards.

Methods: 1006 online participants completed three intertemporal choice tasks followed by a series of self-report current and childhood socioeconomic measures.

Results: An ANCOVA with continuous SES as a covariate on decision preferences revealed a main effect of the reward domain ($F(2,1989) = 88.61, p < .001$) and current SES ($F(1,1989) = 27.50, p < .001$). Replicating previous findings, we find that lower current SES ($r = -0.162, p < 0.001$) was predictive of preferences for immediate monetary rewards. We also find a similar relationship for current SES and health rewards ($r = -0.162, p < 0.001$). However, we did not observe a relationship between SES and temporal discounting of social rewards. Overall, this suggests that the relationships between SES and time preference is stronger for money ($z = -3.71, p < .001$) & health ($z = -3.06, p < .005$) than for social rewards.

Conclusion: Our results demonstrate that individuals from lower socioeconomic backgrounds exhibit a preference for immediacy for health and monetary rewards that is not observed for social rewards. Importantly, this suggests that lower SES is not associated with a universal preference for immediacy, but rather reward-specific. One possible explanation is that SES places specific constraints on monetary and health choices, which shape discounting behavior. Future research can expand upon these findings by examining health discounting in societies outside of the US where health care is universally provided. This can provide further insight into the cultural and systemic factors that shape temporal discounting.

P1-G-27 - Quantifying a task-invariant Bayesian prior for active avoidance: A pilot study

Tobias Granwald¹, Peter Dayan², Máté Lengyel³, Marc Guitart Masip¹

¹Karolinska Institutet, ²Max Planck Institute for Biological Cybernetics & University of Tübingen, ³University of Cambridge & Central European University

In a Bayesian framework, decision-making involves combining information available at the time of the decision with prior experiences or expectations. The prior is especially important since available information can often be ambiguous or incomplete. As such, priors are thought to have a profound impact on behaviour. For example, the generalised passive behaviour in patients suffering from learned helplessness and depression has been conceptualised as resulting from a pessimistic prior implying the expectation of failure when performing an action.

When measuring these priors in the lab, it is important that they generalise across different situations. Otherwise, the priors captured may only reflect task-specific assumptions that do not apply more widely. Task-invariant perceptual priors have been measured using cognitive tasks, however, task-invariant priors for affective decision-making have not been quantified.

To quantify a generalizable prior that actions will be successful, we administered two differently framed decision-making tasks. In the tasks, 54 participants made repeated decisions as to whether to take an active action or not. The passive choice was associated with a sure loss or missing out on a potential reward while the active action was associated with a cost and a probability to avoid the negative outcome. Thereby, the value of the passive choice is fully transparent. We then used computational modeling to characterize each participant's prior expectation that an active choice will result in avoiding a negative outcome. Parameter and model recovery for all models were satisfactory.

We found that a Bayesian prior model fit better than 6 alternative models (Bayesian model with a beta distribution prior vs the best-fit model without Bayesian updating: Task 1, model frequency (MF)=.648, exceedance probability (EP)=.985; Task 2, MF=.658, EP=.990). Furthermore, the means and variances of the priors correlated across tasks (M: $r=.407, p=.003$, var: $r=.271, p=.047$). When explicitly modelled, slightly more participants were fitted by a model with a common prior across the tasks than a model with one prior for each task (MF=.535, EP=.692). Furthermore, the mean of this common prior was moderately reliable one week later ($N=43, ICC(2,1)=.569$). The common prior also correlated with positive mood (M: $r=.244, p=.075$, var: $r=-.310, p=.022$).

In conclusion, our data show that it is possible to use decision-making tasks to quantify prior beliefs that an active choice will result in avoiding a negative outcome. This method will be used to quantify these priors in clinical populations. This may provide insights into learned helplessness and depression.

P1-H-28 - Context-dependent P2 plasticity dynamics in the monetary incentive delay task

Anastasia Grigoreva¹, Alla Kondratenko¹, Aleksei Gorin¹

¹HSE University

Recent advances in neuroeconomics and decision-making research have emphasized the neural mechanisms underlying feedback processing in reinforcement learning. Understanding how humans process feedback, predictions, and prediction errors, as well as their relationship to reinforcement magnitude, is crucial both for elucidating the mechanisms driving human economic behavior and for informing the development of new models of neuroplasticity and learning-related brain reorganization. In reinforcement learning, the value associated with predicted outcomes is also of key importance. It is postulated that reward or loss prediction errors generate adaptive signals that drive behavioral changes.

The aim of our study was to investigate learning-associated plastic changes by examining alterations in evoked potentials in response to auditory cues during one of the most used empirical trials in neuroeconomics – monetary incentive delay (MID) task. To do so, we recorded EEG during the loss version MID task performance and assessed how responses to auditory signals reflected potential

monetary loss evolved. We focused on the P200 evoked activity component, which has been linked to improvements in sound perception and performance in numerous auditory tasks.

We analyzed data from 32 participants who completed the experiment over two days. We compared event-related potentials (ERPs) between the first and second day of the task and across different contexts, anticipating significant differences in P200 amplitude. A variance analysis of the P2 component amplitude revealed the significance of Day factor ($F(1, 31) = 12.711$, $p = 0.00120$, $\hat{p}2 = 0.29$) and Day*Context factors interaction ($F(2, 62) = 4.7554$, $p = 0.01198$, $\hat{p}2 = 0.13$). Post-hoc analysis using the Bonferroni method highlighted significant differences in contexts involving high losses and distinct losses.

Our findings demonstrate that associating a previously neutral auditory stimulus with a specific monetary loss can induce plastic changes, manifesting as context-dependent enhancement of auditory evoked potentials amplitudes. These results contribute to existing models of learning and decision-making and add to our understanding of the neural mechanisms underlying reinforcement learning.

P1-H-29 - Accounting for Multiscale Processing in Adaptive Real-World Decision-Making via the Hippocampus

Dhruv Mehrotra¹, Laurette Dube², Goktug Bender²

¹Integrated Program in Neuroscience, McGill University, Canada, ²McGill University

For adaptive decision-making in diverse and changing real-world contexts, the brain needs to allow past information to influence current processing over multiple timescales, enabling complex and reciprocal dynamics between the two underlying reinforcement learning strategies: model-free (MF) - a reflection of the past - and model-based (MB) - a deliberation on causal environment-behavior structure. Yet, while the role of the hippocampus in memory and spatial learning is well-known, its interactions with the striatum with regards to the “here and now” decision-making is just starting to be explored. This work aims to better appreciate the role of the hippocampus in decision-making. Typically, humans rarely act in a pure MF or MB manner. We therefore review literature in reinforcement learning that examines the successor representation (i.e., an expectation of future states from a given starting state) as a means to bridge the MF and MB strategies. Alongside, we review hippocampal sequences such as hippocampal replay, theta sequences and temporal coding in the hippocampus, demonstrating how the hippocampus can implement MB learning. We argue that the implementation of such sequences in reinforcement learning agents improves their performance, as demonstrated by both successor representation and continual learning agents. The incorporation of sequences enables these agents to perform not only more complex tasks, but also multiscale temporal processing in a biologically plausible manner. Altogether, we articulate an integrative framework to advance current striatal-focused decision making to better account for multiscale mechanisms underlying various real-world time-related concepts such as the self that cumulates over a person's life course.

P1-H-30 - Reward-Driven Superstitions in Uncontrollable Environments

Haoyang Lu¹, Tianyuan Teng¹, Hang Zhang¹

¹Peking University

Objective: Superstitions are false beliefs about causality, often related to an illusory control of luck. Although superstitions are ubiquitous and relevant to our everyday lives, previous research has mainly focused on their historical and cultural background, individual differences, and influencing factors. The cognitive processes underlying the formation and persistence of superstitions have been understudied. In this study, we aimed to quantify and investigate the formation of superstitious beliefs in a freely explorable but uncontrollable environment.

Methods: We designed a behavioral task with different reward contingencies to track the evolution of superstitious beliefs over time. The task involved three between-participant conditions that differed in the mean and variation of reward probability. We recruited 60 participants for each condition. During the task with 100 trials, participants attempted to produce rewardable keypress patterns that were the combination of four keys by themselves. However, unbeknownst to them, the outcomes were independent of their actions. After the task, participants reported their perceived controllability and reward probability during the experiment. They also completed several questionnaires related to their inclination to have superstitious beliefs in real life.

Results: To quantify the evolution of superstitious beliefs, we used multidimensional scaling to map the distances between keypress sequences to a new two-dimensional space. The 340 originally different sequences were then classified into distinct 50 keypress patterns. Similarity analysis of keypress patterns revealed that participants adopted the win-stay, lose-shift strategy ($F(1,179)=227.13$, $\hat{g}2=.11$, $p<.001$), a heuristic to maximize rewards. Those under high reward rates preferred the different sequences within the same pattern after losing and converged to fewer patterns compared with other conditions ($F(2,179)=8.89$, $\hat{g}2=.016$, $p<.001$). Although they had no control of the rewards, participants' perception of controllability was higher under high reward rates, but it increased with their perceived, rather than actual, reward rate (for perceived reward rate: $\hat{r}^2=0.39$, $t(177)=3.72$, $p<.001$; for actual reward rate: $\hat{r}^2=-0.51$, $t(177)=-1.49$, $p=.38$). We also found that the strength of participants' belief in everyday superstitions could predict their reported controllability ($\hat{r}^2=0.18$, $t(168)=2.53$, $p=.012$).

Conclusions: By quantifying the evolution of human behaviors in uncontrollable environments, we find that the formation of superstitious beliefs may be motivated by reward maximization but is associated with a distorted perception of the reward rate.

P1-H-31 - Ambiguity and confirmation bias in reward learning

Hayley Dorfman¹, Rahul Bhui²

¹Harvard University, ²Massachusetts Institute of Technology

Objective: We tend to interpret ambiguous feedback in ways that confirm our preexisting beliefs. We propose that this confirmation bias constitutes an inductive bias that speeds up learning, analogous to missing data imputation. We develop and test this theory using a reinforcement learning (RL) task in which participants are only provided partial information about outcomes, allowing more leeway for subjective interpretation.

Method We introduce a computational model of reward learning from feedback that is sometimes ambiguous. Specifically, we consider valence ambiguity, where people are told the exact magnitude but not whether it is positive or negative. The agent imputes these outcomes with an estimated reward and corresponding uncertainty, according to Bayesian reasoning about the outcome's expected value given the partial information combined with their prior beliefs. To test this model, we developed a behavioral task where participants ($N = 68$, MTurk) chose between two gold mines. After each choice, they saw either a positive (gold), negative (rocks), or ambiguous (dirty sediment) outcome and reported whether they thought the ambiguous outcome was positive or negative.

To manipulate prior beliefs, participants were told at the start of each block that the area they were going to mine in had either an abundance of gold (rich condition), a dearth of gold (poor condition), or that the area was unknown (neutral condition). Participants also completed a measure of trait optimism.

Results The Bayesian model received stronger support according to the protected exceedance probability (PXP = 0.96) compared to other models, though a version with condition-specific priors fit best according to AIC values (three-prior: AIC = 5255, one-prior: AIC = 5366, single learning rate: AIC = 6012). Stated beliefs about ambiguous outcomes varied across conditions ($F(2,134) = 28.82$, $p < 0.001$) and closely reflected true outcome probabilities. This variation seemed best captured by the three-prior model. Participants' average stated beliefs about ambiguous feedback were positively correlated with trait optimism ($r = 0.32$, $p = 0.009$).

Conclusion: We find that people learn from ambiguous feedback by integrating prior beliefs into their value estimates in a confirmatory manner. Subjective beliefs about ambiguous outcomes varied with background information and were related to trait optimism. These results provide a theoretical framework for understanding the effect of valence ambiguity on learning and provide a mechanistic explanation for the rise of confirmation biases.

P1-H-32 - Computational mechanisms underlying latent value updating of unchosen actions

Ido Ben-Artzi¹, Yoav Kessler², Bruno Nicenboim³, Nitzan Shahar¹

¹Tel Aviv University, ²Ben Gurion University of the Negev, ³Tilburg University

Research suggests that humans learn action values through an iterative process of trial and error, whereby the value of action is updated according to observable and temporally adjacent outcomes. However, often the outcome of actions we deliberated upon but did not commit to remains hidden. For example, we might deliberate on which form of transportation is best to commute to our first day at work, by bus, or by train. If we choose to take the bus, the train experience remains unavailable to us, thus limiting our ability to update its value. Reinforcement learning models have extensively examined how humans update action values according to observed outcomes. However, a fundamental question remains regarding whether and how humans may falsely update the value of unchosen actions for which no feedback was observed. 178 participants completed a multi-armed bandit task in which they were asked to make card selections to gain monetary rewards. In each trial, participants were offered two cards (randomly selected by the computer from a deck of four) and then were presented with the outcome for the chosen card. We selectively analyzed trials in which a previously unchosen card was reoffered with a third card that was not presented in the previous trial. We were interested in studying whether the previous outcome for the chosen card would affect the tendency to now choose the previously unchosen card. We found strong evidence for the influence of the previous outcome on participants' choices, such that participants were less likely to choose a previously unchosen card in trial $n+1$ if the chosen card in trial n was rewarded (43%) vs. unrewarded (49%; posterior median = -.25, HDI95% = -.33 to -.17; $pd \sim 100\%$). Contrary to our intuition, we did not find properties of the deliberation to moderate the extent of value updating for unchosen actions. Rather, reinforcement learning computational modeling suggested that for every action, individuals consider the outcome history for choosing an option, together with the outcome history of rejecting the alternative. This mechanism predicted the effect of value update for unchosen actions both on the group and individual levels.

P1-I-33 - Neural representation overlap predicts the decoy attraction effect

Asaf Madar¹, Ido Tavor¹, Dino Levy¹

¹Tel Aviv University

The decoy attraction effect is one of the most well-known effects in behavioral economics. Over the last 40 years studies have replicated and investigated the effect, focusing mainly on structured stimuli with explicit numerical attributes. In the last decade, several attempts have failed to replicate the attraction effect when using more natural stimuli such as images of real-world objects, without providing explicit properties. Nonetheless, the current literature provides no tools to select natural stimuli that would elicit the attraction effect, as the stimuli are not represented by specific well-defined numerical attributes. Here, we propose a novel method for determining which stimuli would elicit the decoy attraction effect, without using any of their explicit attributes, based solely on the similarity of the neural representations of the stimuli as measured by functional magnetic resonance imaging (fMRI). We hypothesize that an overlap in the neural representation between the target and decoy options, but not between the competitor and decoy, would predict the decoy attraction effect. We report preliminary results using lotteries, which are known to elicit the decoy effect and have explicit numerical attributes (amount and probability), but we use only their neural representation in our analysis. First, one group of subjects went through a fMRI scan while performing a Becker-DeGroot-Marschak (BDM) task for each lottery at a time, measuring their willingness to pay for participating in that lottery. Then, for each subject we extracted the neural representation of each lottery, and computed the neural similarity between each pair of lotteries. Using these lotteries, we created lottery sets that elicit the decoy attraction effect. A different group of subjects performed a standard behavioral decoy effect choice task, choosing between the sets of lotteries. Based on this behavioral experiment, we computed the magnitude of decoy effect for each set of lotteries, ranging from 0.1% to 17% in preference change of the target option. Finally, we use the average neural similarity of the lotteries from the BDM-fMRI experiment to predict the magnitude of the decoy effect of each set of lotteries from the behavioral experiment. We find preliminary evidence supporting our initial hypothesis, showing that high neural overlap between the target and decoy corresponds to larger decoy effects, while low overlap corresponds to smaller decoy effects. This suggests a general neural mechanism that underlies the decoy attraction effect which could be used for any natural stimuli, even without any explicit attributes.

P1-I-34 - Improvements in a Decision Support System from the evaluation of cognitive responses with EEG and Eye-Tracking

Anderson Lucas Carneiro De Lima Da Silva¹, Jefferson Souza Medeiros¹, Rayara Sirlei Mauricio De Oliveira¹, André Marques Cavalcanti¹, Ana Paula Costa¹, Lucia Reis Peixoto Roselli², Danielle Moraes¹, Adiel de Almeida¹

¹Universidade Federal de Pernambuco, ²Federal University of Pernambuco

With the advancement of human needs, decisions began to be considered as important in defining the course of society, therefore, multicriteria decision making is a common practice in people and companies. A known method in this context is FITradeoff, which is implemented in a Decision Support System (DSS) (DE ALMEIDA et al., 2021). Due to its flexibility, FITradeoff has been widely used to solve various decision problems. However, there are still unresolved questions about the cognitive demands imposed by the DSS on the decision maker (DM) in the different step of the FITradeoff.

Objective: To carry out a study to evaluate the participants' cognitive behavior when executing the preference elicitation process with

the FITradeoff in order to generate new insights and improvements in the DSS, considering the effect of the number of criteria present in the decision problems.

Methodology: Three experiments were carried out in the NSID Laboratory using an EEG and an Eye-Tracking. The sample consisted of 36 undergraduate and graduate students. In the first experiment, participants were asked to solve mathematical calculations and watch music clips. In the second, the FITradeoff DSS was used to solve a decision problem. From the insights obtained in the two experiments, a third experiment was applied to evaluate the effects of the implemented changes in the DSS. For the analyzes using Eye tracking, heat maps and sweep paths were used. With regard to the EEG, the cerebral potentials of Alpha and Theta and Apha frontal asymmetry were analyzed.

Results: Comparing the brain responses between the first two experiments, it was possible to observe a greater approximation of the global evaluation step and the evaluation by decomposition step with the calculation activities when there are few criteria. When the decision problems had many criteria, only in the evaluation by decomposition there is proximity to the calculation activities. With the results of the eye-tracking, it was verified that the information search behavior remains consistent in the evaluation by decomposition independent of the number of criteria, while in the global evaluation, there is an effort minimization effect when there are many criteria. Thus, in the third experiment, changes were applied in the global evaluation step, incorporating characteristics of the evaluation by decomposition, which resulted in more consistent results.

Conclusion: This research contributes to studies in the field of multicriteria decision making and improvement of DSSs. Its results provide valuable insights for the improvement of the FITradeoff DSS and can be applied to expand knowledge in several other contexts.

P1-I-35 - Decomposing the mechanisms and timing of context effects

Liz Izakson¹, Minhee Yoo², Adam Hakim¹, Ian Krajbich³, Ryan Webb⁴, Dino Levy¹

¹Tel Aviv University, ²Ohio State University, ³University of California, Los Angeles, ⁴University of Toronto

Objective) Would you prefer a Twix bar from a gourmet market or a budget shop? Assimilation and contrast effects predict that the Twix bar from the high-value context should be more or less appealing, respectively. These theories assume that context alters the evaluation of an item. However, context may also change initial expectations without affecting evaluations. Here we examine the effects of context on choice and bids in a design where some options become irrelevant, and assess whether they are due to evaluations or expectations. Methods) We ran two studies: choice and bidding. In the choice study (N = 24), subjects saw three consumer items on each side of the screen. Four seconds later, one item on each side became the target of choice, making the others irrelevant. Targets could be surrounded by high- or low-value context. We modeled the effect of context on the starting point and the drift rate of the drift diffusion model (DDM). In the bidding study (N = 101 Online, N = 28 Lab), subjects saw three items. Four seconds later, one item became the target that subjects were asked to bid on. The target could be surrounded by high-, medium-, or low-value context. We modeled the effect of context on the starting point and the drift rate of a price accumulation model, a continuous-response version of the DDM. Results) Context had a positive effect on choice and bids. Subjects chose the higher-value target more when it was surrounded by high-value context ($\hat{\rho}=0.22$, $p<0.01$) and less when it was surrounded by low-value context ($\hat{\rho}=-0.54$, $p<0.001$). In Study 2, subjects bid more on the target as the value of the context increased (Online: $\hat{\rho}=0.02$, $p<0.001$, Lab: $\hat{\rho}=0.04$, $p<0.001$). However, these positive effects of context were observed only in fast responses, indicating a bias in expectations not evaluations. The modeling supports that conclusion. The DDM analysis showed that context had a significant effect on the starting point (M = 0.03, HDI = [0.01, 0.05]) but an insignificant effect on the drift rate (M = 0.05, HDI = [-0.05, 0.15]). Simulations of the price accumulation model suggest that starting point is the most likely explanation for the context effect that disappears over time. Conclusions) Our results indicate assimilation rather than contrast effects in consumer choice and bidding, primarily due to the influence of initial expectations. Decomposing the choice and bidding processes reveals that context effects occur for the fastest decisions due to starting-point biases. Thus, in a setting in which contextual items become irrelevant later in a decision, context alters decisions primarily through expectations rather than evaluations.

P1-I-36 - Disentangling forms of exploration in a multi-armed bandit task

Nick Sidorenko¹, Hui-Kuan Chung¹, Boris Quednow², Helen Hayward-Könnecke³, Alexander Jetter³, Philippe Tobler¹

¹University of Zürich, ²Psychiatric University Hospital Zurich, ³University Hospital Zurich

In changing environments, exploring lesser known but potentially more rewarding alternatives is adaptive and computational models have proposed to distinguish uncertainty-driven exploration from random behavior. However, these different behavioral patterns are hard to discern in paradigms typically used to study them, such as the dynamic multi-armed bandit task, where reward and uncertainty-related information are confounded. For example, continuous selection of the most rewarding arm automatically makes this arm less informative while increasing the informativeness of less rewarding arms. Moreover, the fact that agents always receive feedback after choosing an arm also complicates distinction between deliberate choices and mistakes. To overcome these issues, we modified the classical task and separated reward from information provision by adding non-informative arms to standard informative arms. In this design, continuous selection of a non-informative highest-payoff arm reflects reward-maximizing behavior, whereas a switch to a non-informative low-payoff arm constitutes a mistake. Assessing both task-related features (model-free analysis) and a computational framework (model-based analysis) in pilot cohort, (N=64), we found that in our task, some choices can be unequivocally categorized as uncertainty-driven exploration, while others can be indeed classified as mistakes, as they are worthless in terms of both reward and information gain. Moreover, we show that some non-random exploratory choices are guided by the objective (but also subjective) value associated with the chosen arm rather than by its uncertainty. Repeated performance of the task (with a gap of two weeks) revealed temporal stability of these patterns (between-session correlation coefficients > 0.45). Next, we replicated our findings in a larger cohort (N=160), while investigating how different neuromodulatory systems contribute to explorative decision-making (between-subject design) in the novel task. Dissociating individual propensity for mistakes from other explorative strategies, we found that cholinergic upregulation specifically reduced mistakes. Collectively, our results highlight the importance of decoupling mistakes from deliberate choice while qualifying and quantifying exploration and suggest a more granular consideration of directed exploration. More generally, our findings causally link the cholinergic system to the resolution of explore-exploit trade-offs in humans and thereby inform the research agenda on therapeutic approaches for disorders involving the cholinergic system such as Alzheimer's or Parkinson's disease, in which such decisions are known to be suboptimal.

P1-K-37 - The association between gaze and choice in value-based decisions may arise post hoc—that is, after the decision

Ariel Zylberberg¹, Michael Shadlen¹

¹Columbia University

Decision makers tend to choose the option they attend to longer. The attentional drift diffusion model (aDDM) is the leading explanation for this observation. It proposes that decisions are made by accumulating noisy evidence over time and that attention, revealed through the gaze, has a causal influence on choice because the value of unattended options is discounted.

Here we consider an alternative explanation for the pattern of observations adduced to support a causal influence of gaze on choice. Response times in simple decision tasks include a decision time and a non-decision time. The decision time encompasses the epoch in which evidence is evaluated, ending in the commitment to a choice. The non-decision times include sensory and motor delays that are independent of decision difficulty. We investigated the hypothesis that the association between gaze and choice emerges after the decision has been made, during the interval between committing to a choice and reporting it.

We evaluated a Post-Decision Gaze (PDG) model, in which attention does not affect the evaluation of evidence. The gaze simply reflects the choice after the decision process has terminated. The PDG model explains key observations that have been taken as support for the aDDM, including the last-fixation bias, the gaze-cascade effect and the effect of the total value of the alternatives on response time. The aDDM and the PDG model make qualitatively different predictions about the link between looking time, response time and choice consistency, which we evaluate with data from the snacks task of Krajčich and colleagues (Nature Neuroscience 2010). Specifically, the aDDM predicts that the time spent looking at the chosen item minus the unchosen item (ΔD_{well}) is greater for choices inconsistent with initial valuation. In contrast, the PDG predicts no difference in ΔD_{well} . The data favor the PDG model. Further, we reanalyze recent proposals suggesting that value-based decisions result from optimal Bayesian inference about the latent value of the alternatives, and show that their fit to the behavioral data can be greatly improved if the association between attention and choice is considered post-decisional. The PDG model does not explain the association between first dwell duration and choice. We discuss potential extensions to the PDG model that might account for this regularity.

Our results challenge widespread assumptions about the role of attention in decision-making and suggest that the apparent causal influence of attention on choice may be at least partially explained by a post-decision mechanism.

P1-K-38 - Effects of encoding precision and comparison difficulty on attentional discounting

Da Li¹, Todd Hare¹

¹University of Zürich

Study objective: Sequential sampling models are used to quantify the key influence of visual attention on decision making. Typically, attentional discounting (AD) is assumed to be constant over time and independent of decision context. However, Brus et al. (2021) found that AD was associated with lower decision confidence and comparison noise. Furthermore, estimates of AD differed between models that did or didn't account for encoding noise. Here, we further examine the relationships between AD and stimulus or decision features.

Methods: We fit the Gaze-weighted Linear Accumulator Model (GLAM; Thomas et al., 2019) to the 24 participants tested in Tavares et al., 2017. In the task, participants were presented with two lines and asked to choose the line with an orientation closer to a target. We estimated 3 parameters, attentional discounting (γ), drift scaling (v), and integration noise (s), in trials of either fine/coarse encoding (EncF and EncC) or easy/hard comparison (CompE and CompH). We calculated 95% HDI of the posterior group mean differences for each parameter across the conditions. For CompE and CompH, the two lines have more distinct or similar orientations, respectively. In terms of encoding precision, previous work on visual perception has shown that lines closer to cardinal orientations are encoded in a more precise, finer grained manner (EncF) compared to diagonal orientations (EncC). We use this well-established feature of human perception to test the role of encoding precision in attention discounting.

Results: We found that encoding precision selectively affects γ (attentional discounting), but not v (drift scaling) nor s (integration noise) in these data. Specifically, γ is higher (discount less) when encoding is more precise (95% HDI for $\gamma(\text{EncF}) - \gamma(\text{EncC}) = [.04, .34]$). Separately, when comparison is easier, γ is lower (discount stronger), v is higher, and s is lower (95% HDIs CompE – CompH: $\gamma = [-.34, -.05]$, $v = [.02, .27]$, and $s = [-.07, -.01]$). There was no significant interaction between encoding precision and comparison difficulty.

Conclusions: We found that attention discounting rates were sensitive to both comparison difficulty and encoding precision. Greater discounting during easy relative to hard trials suggests that discounting may vary as a function of the effort required to perform accurately. The influence of perceptual encoding on the AD is intriguing because it suggests that early sensory processes may feed into downstream interactions between stimulus features and attention. Both results provide further evidence that AD is a dynamic phenomenon rather than a static individual trait.

P1-K-39 - EEG neural correlates of regret and rejoicing: a neuroeconomic experiment

Łukasz Tanajewski¹, Illia Kuznietsov², Marcin Furtak³, Oleksandr Zhuravlov², Martin Schürmann⁴

¹Kozminski University, ²Lesya Ukrainka Volyn National University, ³Nencki Institute of Experimental Biology, ⁴University of Nottingham

Aims: Everyday life requires a train of decisions. Humans may learn about the outcomes of both chosen and unchosen options, and experience regret (when an unchosen option would have been better compared to a chosen one) or rejoicing (when an unchosen option would have been worse than a chosen one). Previous studies have identified the neural signatures of regret, but the regret-rejoicing asymmetries and intensities of these emotions have been poorly explored. To fill the gap, controlling for the reward magnitude of a chosen option and allowing for the outcomes in gain domain only (therefore excluding loss aversion), we studied the neural correlates of regret and rejoicing emotions.

Methods: Forty-five participants made 360 choices between two identical cards when undergoing 64-channel electroencephalography (EEG). The reverse of each card indicated a cash reward of 1, 5, or 9 PLN. Gained reward X on the reverse of a chosen card was revealed for 2000ms and, after a 500ms break, outcome Y on the reverse of an unchosen card was presented for 2000ms. Such trials were pseudorandomized to get 40 trials per condition cXY, where XY = 11, 15, 19; 51, 55, 59; 91, 95, or 99. Event-related potentials (ERPs) were stimulus-locked to revealing an unchosen card, which aimed at studying the neural correlates of regret ($Y > X$) and rejoicing ($Y < X$). We captured the emotions intensities by varying the levels of Y while keeping reward magnitude X (=1, 9) unchanged. Note that all the conditions excluded losses ($X > 0$, $Y > 0$). Based on the literature, the ERP waveforms of interest were: P200, a positive deflection

(latency: ~200ms), and feedback-related negativity (FRN), a subsequent differential deflection (200-300ms), both measured from frontocentral regions; P300, a positive deflection (250-500ms), measured from centroparietal regions. Paired-samples t-test with permutation for multiple comparisons were used ($\alpha = 5\%$).

Results: The P200 amplitude was higher for both c19 and c15 conditions (188-224ms), compared to c11, with mean amplitude for c19 being greater than for c15. FRN was stronger for both c19 and c15, compared to c11, with no difference between c19 and c15. The P200 and P300 amplitudes were lower for both c91 and c95, compared to c99.

Conclusions: We have shown early attention capture in regret (higher P200) and reduced cognitive load in rejoicing (lower P200 and P300). The P200 amplitude could reflect regret and its strength, and so be a neural correlate of regret emotion. These findings suggest that the relevance and valence of unfavorable feedback (depicted by P200), rather than negative prediction error (FRN), are the specific features of regret emotion.

Funding: Polish National Science Center Grant (2017/01/X/HS6/01789). PI: ŁT.

P1-K-40 - The Role of Information Availability in Simple Decisions

Douglas Lee¹, Kostantinos Tsetsos², Giovanni Pezzulo³, Nitzan Shahar¹, Marius Usher¹

¹Tel Aviv University, ²University of Bristol, ³National Research Council of Italy

Most models of choice are based on the idea that people process information about the options before deciding. Recent advances in computational theories have purported that people will seek to process more information about options when they are less certain about the option values. Specifically, they will allocate more (e.g., visual) attention to options with lower certainty, in an effort to increase their feeling of certainty by processing additional information. But such models assume that information is readily available, and that people can process as much as they want whenever they want. I introduce a new variable, information availability, to test what happens when people seek more information but cannot find it. I will present preliminary data from a behavioral / eye-tracking experiment, and I will discuss ideas for how to best model this additional variable.

The objective of this study is to see how stimulus value, uncertainty, and information availability affect choice, RT, confidence, and in particular, gaze patterns. The initial phase of this research is based on a two-alternative forced-choice task between streams of randomly-sampled numbers ($n=50$). The next phase will incorporate computational models, such as the attentional drift-diffusion model (aDDM) and the leaky, competing accumulator (LCA) model. Preliminary results indicate that people gaze longer at options with higher value, higher variability, and lower information availability, and that they choose options with higher value, higher availability, and longer gaze duration. This research could help illuminate the process of information seeking to reduce uncertainty during choice deliberation.

P1-K-41 - Attentional mechanisms underlying cooperative behavior

Claire Lugrin¹, Arkady Kononov², Christian Ruff¹

¹University of Zürich, ²University of Birmingham

Objective: Cooperation is essential for human societies, but not all people cooperate to the same degree. While these individual differences are usually explained by motives such as other-regarding or risk preferences, at least some differences might also relate to how people attend to the choice-relevant information. Here we study with eye-tracking and computational modelling how cooperation is linked to attentional mechanisms.

Methods: 84 subjects played 192 trials of a one-shot prisoner's dilemma (PD) game with an anonymous real opponent while their eye movements were recorded. On each trial, both players chose to either cooperate or defect. While the players' dominant strategy is to always defect, mutual cooperation leads to a higher total payoff than mutual defection. By systematically varying the four payoff values, we created 96 unique PD games, presented in a standard matrix form and randomly varied the position of the options (cooperate or defect) for both the subject (columns) and the opponent (rows) on each trial. Subjects received no feedback except for one randomly selected game paid out at the end of the session. We fit a computational utility model to subjects' choices and used mixed-effects logistic regressions and machine learning models to estimate the interacting effects of game payoffs, information position, and gaze on the subjects' decisions.

Results: As expected, cooperative behavior was influenced by the payoff values ($p < 0.013$) and was linked to other-regarding preferences ($R^2 = 0.74$, $p < 0.001$). However, cooperation was also linked to the relative attention directed to different payoffs: Classifiers trained on gaze sequences were able to accurately predict out-of-sample cooperation (predicted vs measured cooperation rate $R^2 = 0.8$, $p < 0.001$), showing the importance of information sampling mechanisms. Manipulations of the payoff locations significantly affected cooperation rates (row orders $p < 0.001$, interactions of rows and columns orders $p = 0.045$), but not estimated other-regarding or risk preferences ($p > 0.5$). Importantly, the location manipulations altered subjects' gaze behavior and affected the first payoff subjects attended to, exogenously driving their attention to certain gaze sequences favoring higher cooperation rates.

Conclusion: Our results suggest that cooperation does not only depend on payoffs and preferences, but also on attentional mechanisms that can be manipulated exogenously. This has implications for our understanding of individual differences in cooperative behavior and suggests attentional interventions that could enhance cooperation.

P1-L-42 - Good (enough) practices for reproducible neuroeconomics and decision neuroscience research

Ana Martinovici¹

¹Erasmus University

Researchers use a variety of high-frequency physiological data (e.g., fMRI, EEG, eye-tracking, skin-conductance, heart rate) to investigate topics related to neuroeconomics and decision neuroscience. The increasing use of these types of data is a welcome development for science, as it allows researchers to develop theories that could not be tested with behavioral data alone. However, the state of research software needs to improve in order to: (1) make these methods accessible to more researchers, and (2) increase the likelihood of numerical reproducibility for studies that use these types of data.

This poster focuses on tools and practices that improve computational reproducibility of research studies that use high-frequency physiological data. Some of the tools (e.g., version control, automation) are generally applicable and improve reproducibility for a wide range of studies. My examples focus on how to implement these tools while considering challenges specific to studies in neuroeconomics and decision neuroscience research (e.g., large datasets, privacy and data protection requirements, time-intensive

computations). The goal is to: (1) raise awareness and increase adoption of practices that improve numerical reproducibility, and (2) collect feedback from researchers on what other problems should be addressed in order to use a community-driven approach to guide future methodological developments in the area of numerical reproducibility.

Note: If there is an opportunity, I am also happy to do a hands-on workshop on numerical reproducibility.

P1-L-43 - When lie-detectors go overboard: preventing the neural predictor of deception from predicting honest selfishness

Sangil Lee¹, Runxuan Niu², Lusha Zhu², Andrew Kayser³, Ming Hsu¹

¹University of California, Berkeley, ²Peking University, ³University of California, San Francisco

Recent neuroscientific studies have raised hopes for the possibility of neural lie-detectors by showing significant out-of-sample prediction of deceptive behavior. Despite these advances, however, less attention has been paid to longstanding and fundamental questions regarding the validity of biomarkers of deception. In the real world, most lies involve selfish motivations. Therefore, it is critical for a neural predictor of deception to avoid being driven by signals associated with selfishness. We show, through an experimental task ($n = 33$), that neural predictors built naively to predict selfish lies have the unintended consequence that they improperly generalize to selfish but honest behaviors, raising the possibility that the behavioral processes underlying selfishness and deception are inextricably linked. We then develop a novel statistical approach to train a neural predictor with the additional constraint that it must have zero out-of-sample predictive power for nuisance signals. The resulting predictor detects deception significantly above chance but does not improperly generalize to selfishness. These results suggest a path forward in overcoming the longstanding challenges involved in establishing validity of neural predictors of deception and offer a way to formally test the neural dissociability of behavioral processes.

P1-L-44 - Neurally Revealed Preference and Beliefs over Resolving Uncertainty

Stefan Bucher¹, Zih-Yun Yan², Paul Glimcher², Peter Dayan¹

¹Max Planck Institute for Biological Cybernetics & University of Tübingen, ²New York University

Objective: Preferences and beliefs are key ingredients of any choice theory. A core idea of economics is that choice behavior reveals which (if any) preferences and beliefs can rationalize observed behavior. This paper seeks to extend this revealed preference methodology beyond behavior: We expand the theory's domain to include a measurable correlate of neural activity that is well understood to depend on – and hence revealing of – reward expectations. Specifically, we leverage the fact that the BOLD signal in the ventral striatum has been shown to encode reward prediction errors (RPE; Schultz, Dayan, and Montague 1997; Niv et al, 2012), which reflect the difference between the obtained reward and the reward expected under a prior belief.

Methods: We introduce a theory of RPE-revealed preference which makes precise what RPE measurements imply about preferences and/or beliefs. Observing a large RPE, for instance, reveals that the obtained reward was highly valued and/or unexpected. In an fMRI experiment designed to test this theory, participants faced binary lotteries with dynamically resolving ambiguity: A “roulette” wheel indicated the probability of winning a prize (\$50 or \$25) or receiving nothing. The wheel's composition was revealed only gradually to participants, so that they could sequentially update their belief of winning the prize. At a random time during this gradual resolution, we elicited participants' lottery valuation as a probability equivalent: the probability of winning \$50 that would make them indifferent to the lottery under consideration.

Results: Behavioral pilot data ($N=35$) show that, in line with theoretical predictions, participants' valuations closely track the lottery's winning probability ($\beta=1.08$; 95% CI=[1.03, 1.12]) when comparing two lotteries with identical prizes (which amounts to pure belief elicitation). In the presence of ambiguity, participants' revealed beliefs are slightly more pessimistic, indicating ambiguity aversion. When the lotteries' prizes are not identical, participants' elicited probability equivalents reveal their degree of risk aversion. The fMRI portion of the study is currently underway, and will allow us to test whether the preferences and beliefs inferred from measurements of ventral striatal activity are consistent with those revealed by behavior.

Conclusion: If successful, our approach promises a rigorous way to test neuroeconomic theories that make joint predictions about behavior and neural data, as well as empirical constraints on models of hitherto unobservable evolving beliefs.

This research was supported by NSF grant 1948752 to Andrew Caplin and Stefan Bucher.

Poster Session 2

P2-B-1 - Can Responsiveness to Marketing-Induced Calorie Estimation Errors be Linked to Individual- Differences in Metabolic Health?

Shemal Doshi¹

¹INSEAD Business School

Objective: Marketers are often accused of using marketing tactics that nudge consumers to choose high-caloric foods, potentially being a factor driving obesity rates across the world. Recent research found that consumers with extreme obesity were more responsive to different marketing tactics as compared to lean control. Interestingly, they also found that a weight loss intervention in the same consumers led to a reduction in their marketing responsiveness. These findings motivate us to investigate if the marketing responsiveness that leads to obesity, might be mediated by individual differences in the consumer's metabolic state, such as insulin resistance. **Methods:** For the different analyses reported, we used data from two different studies. For study 1, we tested 117 male German participants with varied body composition. The second study consists of the data collected by Cornil et al. 2021. For this study, 117 female French participants were recruited, and they were either lean controls or patients with severe obesity. Across studies, we estimated the amount of calories (per portion) of 12 branded snacks (six brands framed as "healthy"; the other six as "indulgent") used to compute Calorie Estimation Error (CEE, %), which is used as the measure of the responsiveness to marketing tactics. Additionally, different sociodemographic, bodily and metabolic characteristics were measured. **Results:** 1) We replicate the findings of Cornil et. al (2021) in our study (study 1). We find that participants are responsive to marketing framing (overall CEE = 31.1%), which is majorly driven by the overestimation of calories in "indulgent snacks" (30.2%). We also found that CEE is linked to body fat (%) of the participant, consistent with the literature. 2) We run the mediation analysis with a sample merged from two studies, with a log of HOMA-IR (measure of insulin resistance) as a mediator between CEE (%) and body fat (%), while controlling for study effects. We found a significant indirect effect of body fat (%) on CEE (%) ($\hat{\beta}^2 = 0.418$, $p < .01$) through the mediator HOMA-IR. The total effect of body fat (%) on CEE (%) was significant ($\hat{\beta}^2 = 0.64$, $t = 3$, $p < .001$), but with the inclusion of the mediator, the effect of body fat (%) on CEE (%) was not significant ($B = 0.2$, $t = 0.69$, $p = .48$). This shows that insulin resistance fully mediates the link between body fat (%) and CEE (%). **Conclusions:** The findings provide the first evidence of an interaction between food marketing, body composition, and metabolic health linked to obesity.

P2-B-2 - Throwing good effort after bad: Evidence for a sunk-cost effect in cognitive effort-based decision-making

Ziqi Fu¹, Sean Devine¹, Mario Bogdanov^{1,2,3}, Yiguo Wang¹, A. Ross Otto¹

¹McGill University, ²Harvard University, ³McLean Hospital

The "sunk cost effect"—the tendency for people to continue investing resources beyond the point of diminishing returns—is well-characterized in the financial domain of decision-making. Yet, relatively little is known about how it might manifest in other domains. Here, we examine whether past exertion of cognitive effort would also elicit a sunk cost effect.

To answer this question, we recruited 36 participants for an investment task that involved making investment decisions in hypothetical projects, characterized by demand level (low vs high effort) and varying success probabilities (40%, 50%, 60%). If they chose to invest, participants needed to solve arithmetic problems depending on the effort level of the project. For some projects, participants received outcomes immediately (no prior investment condition). For others, they were informed that an additional problem needed to be solved or they would lose their "investment" of previous effort (low or high prior conditions, depending on the initial demand level). Participants were presented with updated effort levels and success probabilities and asked if they wanted to continue the current project. We compared their likelihood of investment decisions between no prior investment trials and second-stage decisions in low- and high-prior investment conditions using logistic mixed-effect regression models.

Our results suggest that participants based investment decisions in the no prior investment trials on the expected value of a project as they invested more often in projects with higher success probability ($B(CI)$: 1.54[1.39, 1.69], $p < .001$) and lower effort levels ($B(CI)$: -1.11 [-1.35, -0.88], $p < .001$). Importantly, and consistent with a sunk-cost effect, we found that participants were far more likely to invest in a project of similar success probability after having previously invested effort into it (low prior investment: $B(CI)$: 1.06[0.83, 1.29], $p < .001$; high prior investment: $B(CI)$: 1.21[0.95, 1.47], $p < .001$), which was validated in an online direct replication ($N = 100$; low prior investment: $B(CI)$: 1.70[1.52, 1.88], $p < .001$; high prior investment: $B(CI)$: 1.82[1.59, 2.04], $p < .001$). In addition, if continuing a project required high effort, participants were more likely to keep investing if they had invested a high versus low level of effort in the initial math problem ($B(CI)$: 0.88[0.35, 1.40], $p < .001$). These results provide an experimental demonstration of a sunk cost effect based purely on cognitive effort and suggest that previous effort imparts a decision bias similar to financial investments previously demonstrated to incur sunk cost effects.

P2-B-3 - Body mass index-dependent shifts along large-scale gradients in human cortical organization explain dietary regulatory success

Rémi Janet¹, Jonathan Smallwood², Cendri Hutcherson³, Hilke Plassmann⁴, Bronte Mckeown², Anita Tusche¹

¹Queen's University, ²Department of Psychology, ³University of Toronto, ⁴INSEAD Business School

Objectives: Individuals differ in their ability to regulate their diets. Why? We propose that the brain might solve the problem of flexible, goal-consistent dietary choices by adopting – and shifting – brain states in a multi-dimensional space of principal dimensions of brain variation (gradients). These gradients (manifolds) were originally determined from the decomposition of resting state data from the Human Connectome Project (Margulies et al., 2016). Large-scale cortical gradients have been suggested to constitute a core organizing axis of the brain and describe an intrinsic coordinate system on its basis. Previous evidence of altered gradient organization due to increased body mass index (BMI) highlights the promise of this framework to study long-term dietary success (Park et al., 2021). Here, we tested whether shifts in brain states along established large-scale gradients predict short-term regulatory success in a laboratory food task. We also examined how body mass moderates this link. **METHODS:** Participants ($N=123$, 27 ± 6 years, 84 f, $BMI < 35$) performed an established fMRI food choice task (Hare et al., 2009). They made choices under a non-constraint (NC) and health-focus condition (HC). We projected task-evoked brain states measured under both choice conditions in a three-dimensional space of established cortical principal gradients (Margulies et al., 2016). The Euclidian distance between individuals' natural and regulatory brain states (NC to HC) captured the shift of neural maps due to dietary health goals. Subject-specific cortical shifts in the gradient space were used to predict individuals' behavioral regulatory success. **Results:** Shifts in task-based brain states between natural and regulatory food choices (NC and HC) predict variance in people's dietary regulated behavior in the laboratory food task ($r^2=0.328$, $p=0.015$). Results revealed a significant interaction effect between the magnitude of these neural shifts in the gradient space and BMI ($\beta=3.12$; $p=0.016$,

CI[0.60, 5.65]). Results of a slope analysis revealed that leaner participants (BMI<25) showed smaller shifts to achieve regulatory success, suggesting that their natural and regulatory dietary brain states are more similar (closer together) in this intrinsic coordinate system of basic cortical hierarchies. **Discussion:** We show that variance in dietary success across people can be understood as changes along neurocognitive functional hierarchies (gradients). Our results indicate how cognitive functions like dietary control might emerge from the cortex through varying similarities along established macroscale patterns of organization.

P2-B-4 - Contributions of General Causality Orientations to Competitive and Cooperative Effort-Based Decision Making

Nicholas Angelides¹, Michael Platt¹

¹University of Pennsylvania

This study utilizes The General Causality Orientations Scale (GCOS; Deci & Ryan, 1985) to seek relationships between causality orientations for motivation and decision-making in two social environments: cooperative and competitive. Here, we will present an extensive, model-free and model-based computational analysis of simulated behavioral data based upon hypotheses regarding cooperative versus competitive social effort-based decision-making. The primary function of this study is to examine how humans' sources of motivation may be related to how they compute tradeoffs between effort costs and monetary rewards as a function of social decision making in cooperative and competitive contexts. Data collection, to be completed this summer, involves a lab-based, computerized task where participants will decide whether to exert a specified, low level of effort or a variable, potentially higher level of effort, accomplished through repeated finger tapping, for an opportunity to enter a lottery for varying monetary rewards. In the behavioral task, participants are shown two possible lotteries. One lottery varies effort costs (10-90% of maximal effort expenditure) and rewards (\$1.50 - \$5). The second lottery is set to 10% effort for \$1. Participants complete 88 trials: 44 against a computerized competitor, and 44 with a computerized teammate. Following the game, participants complete a self-report questionnaire via Qualtrics which includes basic demographic information, questions regarding social preferences, and trait-level measures including the AMI, BIS/BAS, and GCOS. Current and future analyses are completed using the "R" statistical programming language. Data were simulated based on three main hypotheses: 1) effortful choices were most probable when effort costs were low and rewards were high, 2) least probable when effort costs were high and rewards were low, 3) participants will have a higher willingness to expend effort for rewards in competitive, rather than cooperative, environments. On average, participants who are labeled "cooperators" are expected to be associated with the Autonomy orientation of GCOS while "competitors" are expected to be associated with the Controlled orientation of GCOS. We believe that such findings would suggest that intrinsic motivation motivates cooperative behavior, whereas extrinsic motivation motivates competitive behavior.

P2-B-5 - Decision-making on food consumption: A fNIRS study

Victoria Moiseeva¹, Mario Martinez-Saito², Ksenia Panidi¹, Vladimir Kosonogov², Anna Shestakova²

¹Higher School of Economics, ²HSE University

A growing number of studies employ fNIRS — a relatively cheap and less complex method that is specifically well-suited to capture brain activity in surface near brain regions such as PFC— to investigate customers' preference of foods. Brain correlates of willingness to pay (WTP) for foods were first unveiled in the vmPFC and dlPFC. Furthermore, vmPFC is sensitive to the healthiness of foods. Its context-dependent value-sensitivity is controlled by the dlPFC. It remains unclear to which extent the fNIRS and fMRI BOLD measurements can be interchangeably used to tap into the prefrontal function associated with the WTP and decisions about food preferences.

In the first study, 40 healthy participants performed a WTP task in which they had to bid for 30 different sweet and salty snack food items while their BOLD responses were monitored using the 16x12 channel NIRSout system. The 'free-bid' trials were randomly interspersed with the 'forced-bid' trials, where participants made no choice. In the 2nd study, participants were presented with 60 different food items, 30 of them were healthy, while 30 others were non-healthy foods. The participants were presented with the question about their preferences of foods in three contexts: healthiness, tastiness, and neutral. With the self-augmented SPM-fNIRS Matlab toolbox, we performed the analysis HbO data and fit a single linear model with 6 predictors, which corresponded to all combinations of free- vs. forced-bid (choice treatment) with low, middle, and high (price) trials in the first study and 14 predictors, which corresponded to most significant combinations of healthy vs. tasty (priming context) with decision to eat vs. no eat (participant response), and healthy vs. unhealthy (food quality) trials in the second study.

In the first study, free bids provoked a higher level of oxygenation in PFC in comparison to forced bids. This indicates that PFC underpins decision-making, because the perceptual and motor components were excluded.

In the second study, we showed the activation in PFC during decisions to eat in comparison to decisions not to eat. Furthermore, we found larger activation in PFC in response to healthy relative to non-healthy food. Third, we found deactivation of PFC during the decision to eat or not to eat specific foods when attention was directed to the food's taste or healthiness, relative to when no attention was directed to food attributes.

Our results suggest that fNIRS technology can be a useful complement to fMRI and EEG in decision-making experiments engaging cortical brain regions located close to the scalp.

P2-B-6 - Metacognitive Monitoring Compensates for Memory Limitations in Open-Ended Decisions

Zhihao Zhang¹, Andrew Kayser², Ming Hsu³

¹University of Virginia, ²University of California, San Francisco, ³University of California, Berkeley

How decisions are made when there is no pre-defined list of options has long been a question that evades the attention of mainstream decision-making research, with notable exceptions in the literature of consumer research (Alba et al., 1991; Lynch & Srull, 1982). In these choices (e.g., picking a fast food restaurant for lunch), the choice set implicitly encompasses a large number of options (e.g., all fast food chains). However, a smaller choice set, generated primarily through memory retrieval, serves as the basis for the choice (Shocker et al., 1991), creating the possibility of leaving out more preferred options due to retrieval failure, leading to suboptimal decisions (Zhang et al., 2021).

While memory limitations affect open-ended decision quality through their effects on choice set formation, it remains largely unclear whether and how individuals may overcome such limitations and improve the decisions they make. In four pre-registered behavioral experiments (total N = 1,235), we identified metacognitive processes as a key mechanism that may compensate for the constraint that

imperfect memory places on decision quality by guiding consumers to expand their choice sets. More specifically, we found that access to a larger external choice set prompts consumers to switch to options that are higher in subjective valuation than their original memory-based choices ($p < 0.001$ for all categories tested), and that metacognitive monitoring of the potential existence of better options outside of the initial choice set robustly predicts choice switches ($p < 0.005$ for all categories tested). Drawing on the literature in metamemory judgments, a feeling-of-knowing prompt immediately after the initial decision demonstrated a strong association with the participants' subsequent switches when provided in a menu ($p < 0.005$ for all categories tested) – the more strongly they felt that they may have omitted something better in their own choice sets, the more likely they would switch to a different option when a larger choice set was provided later. Lastly, we found that consumers adaptively seek information to expand their choice sets, driven by the same metacognitive mechanisms.

Together, these findings reveal the important role for metacognitive processes in consumers' interactions with large choice sets, especially in open-ended decisions often encountered in the real world. Our work provides a novel perspective on bounded rationality in consumer behavior by demonstrating adaptive awareness of decision errors due to cognitive limitations (Rubinstein, 1998).

P2-D-7 - Computational modeling of a novel choice rodent gambling task.

Claire Hales¹, Kelly Hrelja¹, Catharine Winstanley¹

¹University of British Columbia

Gambling disorder (GD) is an addictive disorder characterized by repeated problematic gambling despite severe negative consequences, often linked to poor decision making and impulsivity. Research into GD has been facilitated by animal research using tasks such as the rodent gambling task (rGT), a translation of the human Iowa Gambling Task (IGT). Like in the IGT, rats must learn to avoid "high-risk, high-reward" options to maximize their profits in the rGT. Previously, an individual rat was trained on either an uncued or cued version of the task. On the cued rGT, sugar pellet rewards are delivered concurrently with salient audiovisual win-paired cues that scale with the size of the reward, mimicking electronic gaming machines. Rats performing the uncued rGT are more optimal, making fewer risky choices than those playing the cued rGT, a finding that has been replicated in humans. In this novel version of the task, on each trial rats were able to choose either a cued trial with win-paired cues, or an uncued trial. Male ($n=32$) and female ($n=32$) rats were trained to stable performance. Behavioral data was analyzed using the drift diffusion model to identify the decision-making processes that underlie cued trial choices, and how this is linked to impulsive responding and choices for optimal or risky options. Matching previous findings, rats that frequently choose cued trials pick the suboptimal high-risk, high-reward options more often. Modeling revealed that along with the evidence accumulation rate, decision starting point was also correlated with choice for the cued lever. This suggests a link between the presence of cues and biased decision-making towards those cues. Non-decision time was the key model parameter explaining differences in optimal vs risky decision making, whilst also being linked to impulsive (premature) responding. Risky rats had shorter non-decision times, whilst in optimal rats, the longer non-decision times were driven by trials which were either followed by premature or risky choices. This demonstrates that processing during choice for cues is not independent from impulsivity and risky decision making. Specifically, factors that contribute to the total response time to choose a cue, but are not directly linked to the evidence accumulation process impact behavior later in a trial. The use of a model framework to link choice for cues with other behavioral measures provides insight into how cues alter cognitive processing to increase suboptimal decision making and impulsivity. This has direct relevance to understanding how cues, such as those used on electronic gaming machines, might negatively contribute to maladaptive gambling in GD.

P2-D-9 - The Effect of Acute Stress on Humanitarian Supplies Management.

Maximilian Burkhardt¹, Felix Nitsch², Stefan Spinler¹, Luk Van Wassenhove²

¹WHU Otto Beisheim School of Management, ²INSEAD Business School

Given their exposed position in the humanitarian operations chain, field workers face remarkably challenging and stressful working and living conditions. This may affect their managerial decision-making, and hence the efficiency of operations. While previous basic research suggests that stress can bias decision-making, concrete evidence in the operations context is lacking. The current project, just published in a top operations journal, addresses this gap.

We model the setting of humanitarian supply chain managers as an adapted newsvendor problem in sample of $N = 154$ German undergraduates. Briefly, in this classic behavioral operations paradigm, a supply manager must stock inventory based solely on the predicted demand distribution. This problem in principle allows for an optimal solution but actual behavior is often sub-optimal. Within each treatment, participants play a 50-period newsvendor game, with the same average demand and economic parameters across all treatments. We experimentally expose participants to either one of two degrees of time pressure (high vs. moderate), or noise combined with or without emotional scenes, or a control condition. Using heartrate and self-reported data, we confirm that these manipulations have different effects on two components of the stress response (negative emotional valence and arousal).

Results showed that moderate time pressure ($b=7.457$, $SE=2.430$, $p=.002$) and exposure to emotional pictures ($b=11.318$, $SE=2.468$, $p<.001$) led to improved newsvendor performance compared to the control condition, operationalized by the closeness of supplies orders to the optimal level. High time pressure, however, had a negative impact on performance ($b=-8.017$, $SE=2.513$, $p=.001$). Further, we observed an inverted-U-shaped relationship between the stress reaction and performance, operationalized via a squared predictor term of heartrate ($b=26.454$, $SE=9.063$, $p=.004$). This means that moderate levels of stress led to improved newsvendor performance, whereas high levels of stress flipped this effect. Using 2 models of behavioral biases in the newsvendor problem, we show that changes in performance are driven specifically by differences in pull-to-center bias (all $p<.001$), i.e., the inclination to anchor orders on mean demand. However, this is not the case for the degree of demand chasing, i.e., the tendency to anchor on prior demand (all $p>.05$).

Conclusion: We provide evidence that stress could affect the efficiency of humanitarian operations. This finding posits a novel argument for addressing the problem of structural stress within organizations.

P2-D-10 - Trial-to-trial fluctuations in risk preference originate from Bayesian inference on noisy neural representations

Gilles De Hollander¹, Marcus Grueschow¹, Christian Ruff¹

¹University of Zürich

Recent theories propose that risk aversion can reflect perceptual biases: Decision-makers have no objective information about payoffs but need to infer them from neural representations (Khaw et al., 2020). Bayesian inference leads to an underestimation of larger payoff

magnitudes – especially for high perceptual noise – and apparent risk aversion. Consistent with these theories, we have shown before that individual differences in risk aversion are predicted by the acuity of neural magnitude representations. However, current theories do neither explain why risk preferences vary across time nor why individuals can be risk-seeking.

Here we close these gaps and show that an individual's momentary preference relates to fluctuations in the fidelity of neural payoff representations and the direction of this relationship depends on choice context. We propose a perceptual model of risky choice that explains this finding and that successfully predicts both relative risk aversion and risk-seeking.

Thirty subjects inside an MR scanner chose between two options that were presented sequentially. Notably, the proportion of risky choices related to both the magnitude and order of the offers (all $p < 0.001$), which cannot be explained with traditional choice models. We developed a perceptual risky choice model where decision makers a) represent offers in working memory with different noise than when they are on the screen and b) can have different prior beliefs about the magnitudes of risky and safe payoffs. Hierarchical Bayesian model fitting confirmed that earlier-presented payoffs were represented more noisily, that some subjects assumed that riskier options have higher payoffs, and that some subject types were risk-seeking.

We used a neural decoding approach to measure how the acuity of parietal payoff representations varied from trial to trial. When neural representations were less accurate, subjects responded less consistently ($p = 0.003$) and their risk preferences were further from risk-neutral behavior ($p < 0.001$). Crucially, whether decreased neural acuity led to increased risk seeking or aversion depended on the decision-maker's prior and the size of the presented offers. Neural acuity was linearly related to trial-to-trial fluctuations of the model's noisiness parameter ($p = 0.0085$).

Our new model of risky choice explains both risk-seeking behavior and order effects, and allows us to show how trial-to-trial variability and bias in risky choice can originate from fluctuating noise in parietal magnitude representations. Risk preferences therefore are not fixed – even at short timescales – and critically depend on brain state and prior beliefs.

P2-D-11 - The Processing of Negative Numbers and its Impact on Subjective Valuation

Hannah Seidler¹, Sebastian Olschewski², Benjamin Scheibehenne¹

¹Karlsruhe Institute of Technology, ²University of Basel & Warwick Business School

The perception and integration of numeric information is a prerequisite for many decisions in our daily life. Research in numeric cognition indicates that positive numbers are represented mentally on a compressed mental number line (CMNL; e.g., Izard & Dehaene, 2008), i.e., a concave mapping of the objective number to the mental representation of that number. The CMNL predicts an underestimation of the mean for sequences of positive numbers. This suggests that for a risky lottery, a subjective certainty equivalent below the expected value (undervaluation) may not only be due to preferences, such as risk aversion, but also due to perceptual biases in the mental representation of numbers. We tested whether the assumption of a CMNL also applies to negative numbers and thus partially explains risk-seeking behavior in the loss domain.

In a series of preregistered experiments (total $N = 638$) we tested these assumptions of the CMNL for positive and negative numbers using a Decision-from-Experience Paradigm: In a fixed-sampling paradigm, participants sequentially sampled 20 numbers drawn from an underlying distribution and were incentivized to either submit an economic valuation or estimate the sequence's mean. The sequences contained either only positive numbers, only negative numbers, or both (i.e., mixed sequences) and varied in mean and standard deviation. For each mean level, we tested the deviation of the participant's response from the sequence's mean against zero using one-sided Bayesian t-tests (Morey et al., 2018). We found underestimation and undervaluation for sequences with positive numbers. For sequences with negative numbers, we found overestimation and overvaluation. The bias for negative and positive numbers was symmetric: a two-sided Bayesian t-test supported the hypothesis that the sum of overestimation and underestimation does not differ from zero. Interestingly, the pattern of the deviation was reversed for mixed sequences: we found overestimation for mixed sequences with positive means and underestimation for mixed sequences with negative means. This hints toward different underlying processes for mixed sequences.

Our findings of underestimation for sequences of positive numbers and an overestimation for sequences of negative numbers support the theory that risk-seeking behavior in the loss domain and risk-averse behavior in the gain domain could partly be explained by biases in number perception and not only by subjective preferences.

P2-E-12 - Internal brain activity and decision making: the effect of pre-stimulus midbrain activity fluctuations on self-control

Jakub Skalbani^{1,2}, Łukasz Tanajewski¹, Marcin Furtak³, Marek Wypych³

¹Kozminski University, ²University of Warsaw, ³Nencki Institute of Experimental Biology

Aims: Recent studies have shown that the spontaneous pre-stimulus fluctuations of brain activity may affect higher-order cognitive processes such as risky decision making, cognitive flexibility, memory encoding, and aesthetic judgments. Some studies suggest that pre-stimulus activity may also be relevant to the mechanisms of value-based decision making requiring self-control. Specifically, a higher pre-trial activity in ventral tegmental area (VTA), i.e., the reward system region involved in self-control, predicted more risk-averse decisions. The aim of this study was to check whether pre-stimulus activity fluctuations in key reward system regions affect food choices that require self-control. Four regions of interest (ROI), i.e., VTA, putamen, nucleus accumbens (NAc) and caudate nucleus were identified based on the literature.

Methods: Participants ($n = 49$) rated the healthiness and tastiness of foods and – based on these ratings – we customized the food pairs for each participant. In the fMRI scanner, the participants made 60 choices between a healthier and a tastier food item twice: after memorizing a seven or one-digit number. Such a memory task ensured that participants had not reflected on a subsequent choice before the foods were presented. Pre-stimulus activity was modeled by the finite impulse response (FIR), a method not biased towards a particular hemodynamic shape, and so optimal for assessing pre-trial signals. Pre-stimulus activity was defined as the mean of the betas across 2 time points (3 seconds) preceding food items presentation. ROI analysis was performed using the repeated measures ANOVA test with two factors: self-control (success, i.e., choosing a healthier over tastier food item, or failure, i.e., choosing a tastier over healthier item) and working memory load (high (HL) or low (LL)).

Results: Our analysis revealed a higher VTA activity prior to the successful compared to failed self-control trials ($p = .04$). Putamen, caudate nucleus and NAc activities did not differ between successful and failed self-control. Pre-stimulus activity was also higher in HL compared to LL trials in VTA ($p = 0.03$), right NAc ($p < .01$), right caudate nucleus ($p = .04$), right ($p = .02$) and left ($p < .01$) putamen.

No interactions were found.

Conclusions: Our findings suggest that a higher pre-stimulus VTA activity may bias food-related decision making towards improved self-control. Value-based decisions may be therefore partly affected by irrelevant internal brain activity, which challenges a traditional view of utility-driven decisions. Spontaneous fluctuations in VTA activity may be one of the mechanisms underlying inconsistencies in human decisions.

Funding: The study was supported by Polish National Science Center Grant 'Cognitive resources and self-control of impulses: a neuroeconomic approach' (No. 2017/26/D/HS6/01157). PI: Łukasz Tanajewski.

P2-E-13 - Behavioral evidence for flexible recency-dependent valuation

Yixin Chen¹, Yutong Li¹, Yeshim Onipede¹, Joseph McGuire¹

¹Boston University

Objective: Episodic temporal information can be relevant to estimating the value of decision prospects. For example, the anticipated value of a carton of leftovers in the fridge—or a food item cached by a foraging animal—can depend on how recently it was placed there. Judgments of recency are thought to engage medial temporal lobe systems, which in turn may interact with reward and valuation systems to guide decisions. Here, we tested the behavioral hypotheses that (1) people can use episodic recency to guide value-based accept/reject decisions, and (2) the directional mapping of recency to value adapts across contexts in which outcomes improve or deteriorate with time.

Methods: Participants viewed a series of unique visual items in interleaved “encounter” and “decision” trials. In decision trials, participants viewed an item they had encountered once before and opted to accept or decline it. Accepting an item yielded a reward or penalty that depended on the temporal lag since its previous encounter. In “rising-value” and “falling-value” environments, item value had a positive or negative association with lag, respectively. Participants (n=31) completed two blocks per environment, with environments cued by different full-screen backgrounds, and received outcome feedback after each decision. A subset of participants (n=12) completed additional blocks with no immediate outcome feedback and repeated the task on a second day of testing.

Results: Per-participant logistic regression analyses modeled accept/reject decisions as a function of log-transformed temporal lag. Among participants who received outcome feedback throughout, estimated slopes were significantly greater in the rising-value environment than in the falling-value environment (Wilcoxon signed-rank $p < 0.001$) and differed from zero in the expected direction in each environment. In feedback-free blocks, slopes likewise differed between environments, both on Day 1 (signed-rank $p < 0.002$) and Day 2 (signed-rank $p < 0.001$), implying contextual cues sufficed to guide context-appropriate remapping of recency to value. The difference in behavior between environments increased with task experience, consistent with experiential learning of temporal contingencies.

Conclusions: The results indicate that human participants can use temporal information in a context-appropriate manner to guide value-based decisions. In future work, the experimental paradigm introduced here can be used to investigate the dynamic re-evaluation of outcomes over arbitrarily long time-scales and can be used in human neuroscience studies to investigate interactions between memory and valuation systems.

P2-E-14 - Dissociable sensitivity to effort and delay costs underlies apathy and impulsivity in Huntington's disease

Lee-Anne Morris¹, Kyla-Louise Horne², Laura Paermentier², Christina Buchanan³, Michael Macaskill², Daniel Myall², Sanjay Manohar⁴, Matthew Apps⁵, Masud Husain⁴, Richard Roxburgh⁶, Tim Anderson², Campbell Le Heron²

¹University of Otago, Christchurch, ²New Zealand Brain Research Institute, ³Auckland City Hospital, ⁴University of Oxford,

⁵University of Birmingham, ⁶University of Auckland

Background and objectives: Apathy and impulsivity are pervasive and debilitating behavioural changes that often co-occur in people with Huntington's disease (HD), a genetic neurodegenerative disorder. Alongside disruptions to basal-ganglia-frontal circuitry, the younger age of people affected compared to most neurodegenerative conditions provides an important model in which to study breakdown in the cognitive processes underpinning these behaviours. In the present work we aimed to identify the behavioural signatures of apathy and impulsivity in Huntington's disease, using two effort-based decision-making paradigms: binary choice and foraging.

Methods: Carriers of the Huntingtin gene mutation (n = 53; premanifest to mild motor manifest disease stage) from three movement disorder clinics (Christchurch & Auckland, New Zealand and Oxford, UK) and healthy controls (n = 19) performed the Apple Gathering Task, in which participants made sequential decisions to accept or reject offers of reward (apples with monetary value) for physical effort (squeezing a hand-held dynamometer), across pseudorandomised combinations of reward and effort. A subset (HD, n = 36; controls, n = 16) performed a foraging task, in which physical effort and delay costs associated with patch-leaving were varied. Measures of mood, behaviour, cognition and disease severity were obtained. Choice data was analysed with generalised linear mixed effects models.

Results: Apathy, but not impulsivity, predicted acceptance of fewer offers overall on the Apple Gathering Task, driven by altered responding to changing effort – but not reward – in the apathy group. As effort levels increased, apathetic people were less likely to accept offers (apathy*effort, $p < 0.001$). In contrast, in the patch leaving task, despite also exerting effort after every decision, apathetic people were less sensitive to increasing effort costs to move on from a patch ($p = 0.002$). In this task impulsivity was not associated with changing effort costs, but an increased change in leaving time as delay costs increased (impulsivity*delay, $p = 0.001$).

Conclusions: Apathy in Huntington's disease is characterised by increased sensitivity to effort in a binary effort-based choice task, but reduced sensitivity to effort costs in a foraging context. In contrast, impulsive people were more sensitive to changing time costs (and thus background reward) to move between patches. Despite their common co-occurrence in HD, apathy and impulsivity have dissociable behavioural signatures, pointing to disruption of multiple components of decision-making, as well as the key role of decision context in the computation of costs.

P2-E-15 - Validating a behavioral measure of individual differences in willingness to wait for delayed rewards

Lena Schaefer¹, Yixin Chen¹, Daniel Fulford¹, Joseph McGuire¹

¹Boston University

Objective: Willingness to wait for delayed rewards differs across individuals and across situations. In studies of foraging-style laboratory tasks with temporally uncertain rewards, human decision makers adjust their persistence appropriately for the distribution of delays they experience. The group-average tendency toward context-appropriate calibration is accompanied by substantial individual variation, but contributors to such variation are unknown. Here, we examined the test-retest reliability of a task-based index of persistence and its associations with theoretically related constructs such as inhibitory control, temporal discounting, and self-reported impulsivity.

Methods: We recruited 160 participants (121 female and 39 male, mean age 25.1, range 18-65) for two in-person laboratory sessions 2 to 4 weeks apart. Each session involved two ten-minute blocks of a willingness-to-wait task, with reward-timing statistics favoring limited persistence in the first block and higher persistence in the second block. We used survival analysis to estimate each individual's average giving-up time, accounting for the censoring of observed waiting times by reward delivery events. Participants completed a battery of other measures including a stop-signal task, a questionnaire-based assessment of temporal discounting, and self-report measures of impulsivity (BIS-11 and UPPS-P).

Results: Participant-specific giving-up times were correlated between the high-persistence and limited-persistence environments ($r=0.85$) and between separate days of testing ($r=0.83$). We observed no evidence that giving-up times were associated with a measure of inhibitory control derived from the stop-signal task, which itself showed moderate test-retest reliability ($r=0.50$), or with temporal discount rates. Shorter behavioral giving-up times were associated with higher scores on certain dimensions of self-reported impulsivity, including Attentional Impulsiveness in the BIS-11 ($r=-0.19$, $p=.017$) and Lack of Perseverance in the UPPS-P ($r=-0.32$, $p<.001$).

Conclusion: Our results inform ongoing debates about the relationship between task- and questionnaire-based assessments of individual differences and about the role of control processes in temporally extended decisions. Our behavioral index of persistence showed high test-retest reliability, smaller associations with self-report constructs, and no evidence of a link with inhibitory control. The findings shed light on the dimensional structure of impulsive behavior and establish groundwork for research examining individual differences in persistence in terms of behavior-derived computational parameters.

P2-F-16 - The computational basis of private and social information aggregation in social networks

Wenning Deng¹, Antonio Rangel¹

¹California Institute of Technology

Objectives: Most decisions have to be made with imperfect and noisy information. Examples include whether to buy an insurance, or how to respond to a date invitation. We often use a combination of personal and social information to inform these decisions. Personal information (called private signals in economics) reflects our individual assessment of the situation, based on the information that we have encountered. Social information comes from observing the choices that others make in the same situation. Optimal decision-making requires balancing between these two sources of noisy and incomplete information. This study is designed to characterize the computational basis of this types of decisions at the individual level, as well as the social dynamics that they generate.

Methods: We collected pilot data from 150 subjects, divided into 6 groups of 25 people each. We build our task from classical economics experiments on information cascades. Each round, subjects were randomly placed in a symmetric circulant social network with 4 neighbors each. Every trial, the computer selected a binary true state of the world with equal probability. Subjects were asked to guess the true state based on two pieces of information: 1) a very noisy but informative private signal, and 2) a live update of the current choices of their neighbors. Each trial lasted for 12 s, participants could change their decisions multiple times, and they were incentivized both for correct and fast responses.

Results: We found that people use both personal and social information to make decisions. Participants' decisions are more accurate when 1) their private information was indicative of the truth, and 2) their neighbors' decisions were more accurate. We also found that people's accuracy was boosted significantly through this information cascade, and the effect is larger when early social information is more accurate. We are currently developing and testing sequential integration models that describe the dynamics of individual decision making as well as the dynamics in the group.

Conclusions: Our preliminary results show that sequential integration models provide a reasonably description of both the dynamics of the individual decision processes, as well as the information cascades in the group. Interestingly, the experiments and models make predictions about the information aggregation quality in different locations and time points in the network, which could serve as an entry point for improving decision making in social networks.

P2-F-17 - Does Androstadienone Affect Social Distance-Dependent Prosocial Behavior? A Double-blind, Placebo-controlled Study

Junsong Lu¹, Yin Wu¹

¹The Hong Kong Polytechnic University

In navigating the complexities of social life, humans have evolved to interpret the unseen language of odorous chemical cues, with profound behavioral impacts often unbeknownst to the conscious mind. The manifestation of this in humans is potentially evident in the scent of androstadienone (androsta-4, 16, -dien-3-one), an odorous compound that has been considered as a putative human pheromone. The current study aimed to investigate the effect of androstadienone on social distance-dependent prosocial behavior measured by a social discounting task, wherein participants made choices between selfish and generous alternatives. Based on our preregistration, we hypothesized that the effect would be sex-specific, with males exposed to androstadienone exhibiting increased generosity, while females would choose more selfish options. Employing a double-blind, placebo-controlled, between-subject design, we recruited 170 participants who were randomly assigned to either the androstadienone or control condition. Olfactory stimuli were administered through the nose while participants completed the social discounting task. Inconsistent with our hypothesis, the results revealed that inhaling androstadienone did not impact social distance-dependent prosocial behavior. This finding was supported by multiple estimates of prosociality, including model-free, model-based, and maximum likelihood estimation. Further analyses indicated that androstadienone administration did not influence perceived social distance or bias participants towards either the generous or selfish option in the social discounting task. We discussed potential explanations for the inconsistency between our hypothesis and empirical findings.

P2-F-18 - Psychopharmacology of Social Reward Revaluation

Tobias Kalenscher¹, Sandra Schäßle¹, Irina Noguer-Calabús¹

¹Heinrich Heine University Düsseldorf

Rats adapt their food choices to conform to their conspecifics' dietary preferences, a phenomenon termed social transmission of food preferences (STFP). STFP can be used to study the neuroanatomical and psychopharmacological underpinnings of socially transmitted overriding of endogenous food preferences. We trained observer rats to reveal their endogenous preference for one out of two appetitive food options. Afterward, they were exposed to a demonstrator rat who was fed with the observer's originally nonpreferred food, and the observer's food choices were sampled again. Observer rats typically change their food preferences following interaction with the demonstrator, specifically by increasing the intake of their originally non-preferred food type. We could recently show that the integrity of rat nucleus accumbens shell is necessary for STFP. Here, I will present the results of a number of follow-up studies on the psychopharmacology of STFP. Specifically, I will report findings after selective dopamine lesions in nucleus accumbens using local 6-hydroxydopamine injections, aberrant dopamine processing in DISC1-transgenic rats, systemic fluoxetine (a selective serotonin reuptake inhibitor) administration, local injections of 5,7 dihydroxytryptamine (producing selective serotonin lesion) into amygdala, systemic oxytocin administration as well as behavioural results on an ingroup-outgroup modulation of STFP (the demonstrator was either a cage mate or not). The pattern of results helps to understand the psychopharmacology of conformity – the alignment of one's preferences to those of other individuals.

P2-F-19 - Loss framing increases sensitivity to productivity and luck in distributive decision

Yoonseo Zoh¹, Ryan Carlson², Steve Chang², Molly Crockett¹, Yeon Soon Shin¹

¹Princeton University, ²Yale University

Distributive justice is an essential component of human cooperation. Productivity that is generated by an individual's deliberate effort and ability and non-controllable external conditions, such as brute luck, represent two important factors for decisions to distribute scarce resources. However, it remains unclear whether and how people consider those factors when making resource distribution decisions, and how such decisions are affected by choice framing. Here, we developed a new paradigm in which participants allocated money between two other people on the basis of varying amounts of productivity and luck. In our two pre-registered experiments of $N = 300$ participants, we found that sensitivity to productivity and luck in resource distribution decisions depends on how the consequence of luck is presented. Specifically, presenting the consequence of luck in terms of lost outcome (versus total outcome gained) increased the weight on both productivity and luck in distributive decision-making, leading to increased compensation for equity. Moreover, the relative sensitivity to luck compared to productivity negatively predicted an individual's attitudes toward political conservatism, neo-liberalism, and protestant work ethics. Our finding shows that individuals' valuation of factors contributing to distributive justice permeate their beliefs about social and economic issues but also can be flexibly adjusted based on which specific feature of an external factor is emphasized.

P2-F-20 - Children's Subjective Socioeconomic Status and its Effect on Sharing Behavior in an Online Dictator Game

Sylvia Pinheiro¹, Lisa Bas¹, Anita Tusche¹, Valerie Kuhlmeier¹

¹Queen's University

Though past research has examined the association between objective socioeconomic status and sharing behavior in both children and adults, results have been mixed. Yet, the broader research question is timely and important: the norms for sharing within a child's socioeconomic context are thought to be learned during middle childhood, leading to the perpetuation of social systems of greater or lesser inequity. Here, we take a novel approach to this question: we examine how children's perception of their social status (subjective socioeconomic status, SSS) may relate to their sharing decisions. We consider (1) when children's SSS comes to match that of their caregivers, and (2) how SSS relates to sharing behavior in middle childhood.

In this online study, 253 children aged 3.5- to 11-years indicated their SSS using an adapted 5-point version of the MacArthur ladder task. Caregivers responded using the original 10-point scale. We examined relationship between SSS and sharing behavior of a subset of 169 children aged 6 to 11 years in a two-trial Dictator Game (with ten coins) conducted in either a monitored or an unmonitored interaction. Sharing behavior was considered in terms of number of coins given across 2 trials (0-20) and equitable sharing (giving 10 coins or more across 2 trials).

In line with previous findings, our results of a hierarchical linear regression ($F(2,189) = 28.82, p < .001, R^2 = .23$) show a decline in children's SSS with age ($B = -.38, p < .001$), and a positive association with the caregiver's SSS ($B = .17, p < .01$). The similarity between children and caregiver's SSS is also predicted by age ($B = -.38, p < .001$), with children over 7 years being more likely to be accurate or underestimate their status. Our analysis testing generalized linear models demonstrated that, above and beyond the effects of age and monitoring the study, children in the low- and mid-SSS categories shared a higher number of coins across the two trials of the DG ($X^2(1,169) = 767.52, R^2N = .54; Blow = .65, plow < .001$ and $Bmiddle = .56, pmiddle < .001$), and children in the low-SSS category were more likely to be equitable ($X^2(1,169) = 209.36, R^2N = .23, Blow = 1.26, plow < .05$).

These findings suggest an early onset of consideration of social position, which, in turn, aligns more closely with adult levels during middle childhood. This subjective socioeconomic status appears to influence children's sharing decisions, as lower-SSS was associated with increased coin sharing and more equitable choices. This observation additionally highlights the value of incorporating subjective perceptions of status in the study of pro-sociality.

P2-G-21 - Exploration in decision-making under uncertainty may be regulated by learning, not by offline tDCS

Victor Timokhov¹, Oksana Zinchenko¹, Anna Shestakova¹, Victoria Moiseeva², Vasiliy Kliucharev³

¹HSE University, ²Higher School of Economics, ³Higher School of Economics - National Research University

Decision-making under uncertainty often requires arbitrating between exploiting the current behavioral strategy with some possible adjustments versus exploring new behavioral strategies instead ("exploitation-exploration" dilemma). Koechlin et al. (2014) proposed a neurocomputational model of this process. Specifically, the right frontopolar cortex (rFPC) is assumed to infer the absolute reliability of alternative behavioral strategies (i.e. that are monitored, but not guiding ongoing behaviour) (Donoso et al., 2014). Thus, the aim of this research was to study the causal role of the rFPC in processing and storing information about alternative strategies. We hypothesized that the inhibition of the rFPC with cathodal tDCS will decrease the efficacy of the alternative strategies retrieval.

In total, 43 participants took part in our study ($F = 24$, mean \pm SEM age = 21.1 ± 2.4 , range = 18–29). High-definition tDCS protocol was simulated in SimNIBS and included four anodes over F8, AF8, Fp2, and FpZ with one cathode above the right eyebrow (15 minutes, 1.5 mA). Additionally, we used fNIRS to check physiological effects of tDCS. The goal of the experimental paradigm was to find correct combination of 4 keys and 3 numbers. There were only 3 unique combinations that switch without prior notification, feedback was sometimes unreliable. The procedure included two 5-minute resting state fNIRS recordings before and after stimulation followed with the behavioral task and one 5-minute recording after the task. Participants received sham and cathodal stimulation with a one-week break. Behavioral data was fitted to the PROBE model (Collins & Koechlin, 2012), which classified trials into exploratory and exploitative.

Preliminary results suggest no effect of stimulation ($t(41) = -0.04$, $p = 0.97$) on the number of exploratory trials across all 3 combinations. However, there is a weak statistical trend for the learning effect on the exploration, as on the second day there were slightly less exploratory trials on average ($t(41) = 1.72$, $p = 0.09$). Moreover, noise parameter ϵ , that scales lapses probability in action selection, also showed a statistical trend to be lower at the second day ($t(41) = 1.86$, $p = 0.07$). Importantly, the results might change after accounting for the individual differences between mainly exploring and exploiting participants. Nonetheless, it can be assumed that the learning effect in this study was stronger than the stimulation effect.

Funding: This work is an output of a research project implemented as part of the Basic Research Program at the National Research University Higher School of Economics (HSE University).

P2-G-22 - Cognitive Components of Trust

Xiaoyan Wu¹, Niklas Brger², Christian Ruff³

¹Beijing Normal University, ²University of Zurich, ³University of Zrich

Objective: Trust is essential for all economic and interpersonal interactions. Studies have shown that people base their trust on inferences about others' trustworthiness, but existing studies on this topic have mainly focused on single aspects of social information and have rarely provided explicit models of how these are incorporated in trust inferences. Here we present a new experimental paradigm and modeling framework that allows us to systematically quantify and compare the impact of three different types of social information on trust evaluations and trust-related decisions.

Methods: Fifty-four investors and 80 trustees played a modified trust game preceded by various rounds of laboratory interactions. The investors first had to complete various political and personal questionnaires that allowed us to build real social profiles of them. Subsequently, the trustees were shown pairs of the investors' social profiles and had to allocate varying amounts of money between them. All payoffs were real, and allocations depended systematically on the match between the trustees' and the investors' personal attitudes. In the subsequent trust game, the investors were shown in each round information about the trustee's generosity (how much they shared with both investors), favoritism (how much they allocated to the investor), and reward (the final amount the investor received). In various classes of computational models, we investigated how the resulting trust decisions (with real monetary consequences) depended on these three types of social information, over and above general social preferences and betrayal aversion.

Results: Participants increased their investment with larger rewards (linear mixed-effects model, $p < 0.001$) and generosity ($p < 0.001$). An interaction of generosity \times favoritism ($p < 0.001$) indicated specific trust if generous trustees previously favored the investor. Computational modeling showed that participants in general use all three types of social information to infer trustworthiness, but to varying degrees, even when accounting for individual degrees of social preferences and betrayal aversion. We are currently recording functional magnetic resonance imaging of the investors to identify whether distinct neural mechanisms underlie these different components of trust estimation.

Conclusion: People systematically evaluate and infer the social characteristics of others from previous interactions and use this information to guide their future behavior. Our paradigm and model extend the understanding of the cognitive components underlying trust and may inform interventions aimed at fostering trust in interpersonal interactions.

P2-H-23 - How to behave variable: meta-learning learning rate or meta-learning epsilon?

Janne Reynders¹, Tom Verguts¹, Senne Braem¹

¹Ghent University

Certain real-life situations require us to be variable or unpredictable. Think about sports such as tennis or boxing, where it is beneficial to hit the target in a variable way such that opponents can't predict the next moves. Behavioral experiments in which variable responding is rewarded show that humans and non-human animals are able to increase the variability in their responses strategically. However, it is not yet clear how such variable behavior is established. One hypothesis is that we have an inborn random generator. This is also implicitly assumed in many computational models of decision-making, by relying on parameters that represent the amount of randomness that is used for decision-making, such as epsilon in an epsilon-greedy approach. Others hypothesized that variable behavior results from a dynamic process between reinforcement and extinction, and hence, a high learning rate. To assess these hypotheses, we developed two versions of a Rescorla-Wagner model with an epsilon-greedy decision policy, that either meta-learned epsilon or learning rate using a policy-gradient method. We tested these models in an 8-armed bandit task where reward distributions were either stable, volatile or hypervolatile (each requiring an increasing amount of variable responding). This last context has been used previously to document arguably random behavior, where only the previously least frequently chosen options were rewarded. Both meta-learning (ML) models were simulated 900 times across the three contexts while varying the initial value of epsilon or learning rate. The ML model of epsilon outperformed models in which epsilon and learning rate were kept constant and showed convergence to an epsilon value of 0.81 ± 0.18 in the hypervolatile context (compared to a stable context, 0.039 ± 0.14 , or volatile context, 0.28 ± 0.3). The ML model of learning rate showed a convergence of learning rate to 0.88 ± 0.13 in the hypervolatile context (compared to a stable context, 0.20 ± 0.28 , or a volatile context, 0.73 ± 0.29). Interestingly, the ML epsilon model obtained a similar reward rate as the ML learning rate model in the hypervolatile context, suggesting that both strategies could lead to a high reward rate. Therefore, we also fitted the model on human data in which participants were subjected to a similar hypervolatile context ($n = 6$, 17255 trials on average). There, the epsilon estimations favored the high epsilon model as all six participants showed an average epsilon of between 0.98 and 0.99. Together, our results support the idea that subjects used a maximal amount of randomness, as predicted by the random generator hypothesis.

P2-H-24 - Learning non-linear rewarding rules can be difficult, even in a minimalistic environment

Jianbo Chen¹, Tianyuan Teng¹, Hang Zhang¹

¹Peking University

Objective: People are known for their ability to learn probabilistic rewarding rules of the environment. However, the rules tested in most studies consisted of one feature dimension or a linear combination of multiple features. Here we investigated how well people can learn non-linear combinatory rules where each single dimension is noninformative.

Methods: In an adapted matching-pennies game, we asked participants to predict whether their computer opponent would choose head or tail. On each block, the computer's choice either followed a probabilistic one-dimensional rule, where the opponent had a 0.8 probability to repeat (or reverse) its choice on the previous trial, or followed a probabilistic XOR rule, where the opponent had a 0.8 probability to choose head (or tail) if the two players had matching pennies and tail (or head) otherwise. Participants' performance was measured by the proportion of trials where they selected the side that the opponent was more likely to choose ("proportion of optimality"). We had run a series of experiments (total N=223) to test whether participants could learn the XOR rules equally well as the one-dimensional rules.

Results: We found striking failures in human learning the probabilistic XOR rules. In Experiment 1, participants' proportion of optimality was lower in the XOR-rule blocks (.52, 95% CI [.51, .53]) than in the one-dimensional repeat-rule (.82, 95% CI [.80, .84]) or reverse-rule (.62, 95% CI [.66, .69]) block. To exclude the possibility that the failure to learn XOR was due to the consideration of a potentially vast feature space (e.g., the two players' full choice history n-trial back), in Experiment 2 we explicitly instructed participants to pay attention only to the two players' choices in the previous trial (visible during their decisions), where the results were the same, with the proportion of optimality lower for XOR (.54, 95% CI [.52, .55]) than for repeat (.69, 95% CI [.65, .72]) or reverse (.63, 95% CI [.60, .67]). In Experiment 3, we extended the block length from 80 trials to 320 trials but the proportion of optimality for XOR was still relatively low even in the last 80 trials (.62, 95% CI [.55, .69]). In Experiment 4, we switched to a cover story that encourages the association of the two dimensions, but the results hardly changed. Further modeling analysis shows that people who fail learning XOR rules might rarely explore the rules, instead of rejecting the rules after testing them.

Conclusions: That people fail to learn probabilistic XOR rules even in a small feature space is surprising, which poses challenges to theories of reward learning.

P2-H-25 - Separable neurocomputational mechanisms underlying multisensory learning

Saurabh Bedi¹, Ella Casimiro¹, Gilles De Hollander¹, Arkady Konovalov², Fritjof Helmchen¹, Silvia Brem¹, Christian Ruff¹

¹University of Zürich, ²University of Birmingham

Multisensory learning is crucial for acquiring cognitive skills such as reading, and disruptions of this process have been linked to several learning disabilities. Despite its importance, the neurocomputational mechanisms underlying multisensory learning remain poorly understood - in particular, to what degree multisensory learning is based on reinforcement or detection of statistical structure. Here we investigate these mechanisms with computational modelling and fMRI.

We designed a novel multisensory learning task that allows us to computationally separate reinforcement (RL) and statistical learning (SL) through feedback and statistical regularities in stimulus presentation, respectively. Participants (N=58) completed a task requiring them to learn combinations of visually-presented insect species and matched mating calls (auditory tones) or dances (tactile finger stimulation). For each participant, a Dirichlet multinomial model of the assigned stimulus sequence yielded information-theoretic measures of the surprise associated with presentation of each multisensory combination, while probabilistic reward and comprehensive choice modelling (comparison of 31 competing RL models) made it possible to compute best-fitting reward prediction error signals. Importantly, both types of learning signals varied independently of one another across time ($r=.0531$).

We find behavioral and neural evidence for both types of learning. Even though participants are not explicitly rewarded for learning the statistical regularities, a mixed-effects regression model shows a significant effect of trial-by-trial statistical surprise on response time ($\beta=.03$, $p < .001$). Moreover, people consistently use the information from (missed) rewards to update their beliefs about the multisensory combinations, in a manner that varies systematically with the current level of uncertainty.

Whole-brain univariate analyses of the fMRI data (cluster-corr. FWE $< .05$) show that both RL and SL are reflected by separate neural learning signals: Surprise is encoded in the TPJ and precuneus, while RPEs relate to BOLD signals in reward-processing but also visual and parietal regions. Current analyses focus on separating modality-specific learning signals with multi-variate decoding, as well as on identifying functional connectivity patterns related to both types of learning.

Taken together, our study demonstrates that multisensory learning relies on both statistical and reinforcement learning, and showcases how these two types can be separated, quantified, and related to behavior. This approach may help to better understand the neural origins of learning disabilities.

P2-H-26 - The impact of presentation order and response labeling on the reward learning of task representations

Mina Habibi¹, Mehdi Senoussi², Pieter Verbeke¹, Senne Braem¹

¹Ghent University, ²Université de Toulouse

Humans are efficient at learning different task structures to exploit task-specific reward relations. While many have studied the impact of the complexity and volatility of reward relations, few have investigated what aspects of visual presentation can facilitate or impede the learning of these task representations. These visual aspects may impact the dimensionality of task representations, where high-dimensional representations are thought to increase separability and avoid interference, while low-dimensional representations allow faster generalization but low separability. In this study, we investigated how two aspects of visually presenting nominally identical task representations (in terms of S-R rules and complexity) can influence the reward learning and geometry of task representations. Specifically, we examined how presenting stimulus versus task information first and using same versus different response labels across tasks affected reward learning and the dimensionality of task representations. We ran neural network simulations and conducted behavioral data collection ($n = 228$) where participants learned to classify images of stimuli depending on three different contexts (i.e., tasks) through reward-based learning in four different experimental groups: Context-first groups with and without Response labeling, and a Stimulus-first groups with and without Response labeling. After 288 trials of training, participants entered a testing phase where they were still rewarded for optimal responding, with no feedback. Our results showed that the accuracy decreased significantly

from training to test ($p < .001$), and this decrease interacted with the factor Response labeling ($p < .001$), but not Presentation order ($p = .331$). Providing different Response labels across tasks helped participants retain the learned task structures from training to test. This further interacted with presentation order ($p = .040$), with this effect being more pronounced in the stimulus-first presentation group. Neural network simulations revealed that training with different response labels induced more high-dimensional task representations, suggesting that response labeling can lead to more high-dimensional task representations improving reinforcement learning in a multi-task environment. We suggest that using context- or task-specific response labels can help participants in the reward learning of more high-dimensional task representations, which can be retained better in periods without reward information. We are currently exploring the potential impact of this learned task dimensionality when repurposing these stimuli in subsequent value-based decision-making.

P2-I-27 - System neglect is associated with network-level parameter selectivity

Mu-Chen Wang¹, George Wu, Shih-Wei Wu¹

¹National Yang Ming Chiao Tung University

System neglect – a tendency to respond primarily to the signals and secondarily to the system that produces the signals – underlines a wide array of human judgment biases, from demand forecasting, pricing decisions, to regime-shift detection. The neural mechanisms that give rise to system neglect, however, remain elusive. To address this question, 30 human subjects performed an fMRI regime-shift detection task that reliably elicits system neglect. In each trial, a series of signals were presented to the subjects. These signals, red or blue balls, were drawn sequentially from one of two urns. At any point during a trial, the urn may switch from one to another. The subjects' task was to report their probability estimates that the regime (the urn) has shifted after observing each signal. Crucially, we manipulated the transition probability (how likely the shift happens) and signal diagnosticity (level of noise in the signals). We found that the subjects showed system neglect in regime-shift estimation. Specifically, system neglect led to over- and underreactions to change: Overreactions were most commonly seen when the signals are noisy but when the environments are stable where change is possible but unlikely. By contrast, underreactions were observed when the signals are precise but when the environments are unstable. We found that system neglect arises from dissociable patterns of selectivity in two distinct brain networks. A frontoparietal network selectively represented individuals' degree of neglecting signal diagnosticity but not the transition probability, while the ventromedial prefrontal cortex showed the opposite pattern. Given that system neglect exists in many different domains of judgments, these results suggest that network-level parameter selectivity is an important building block for human judgment biases.

P2-I-28 - Characterizing the causes, dynamics, and consequences of choice deferral

Yi-Hsin Su¹, Xiamin Leng¹, Amitai Shenhav¹

¹Brown University

Objective: In spite of its ubiquity and significant costs, choice deferral remains poorly understood. Past work has shown that how similar and how low in value one's options are both predict the likelihood with which a person will avoid choosing entirely, but whether these factors play a similar role in determining decisions to delay a choice that eventually needs to be made has yet to be investigated. Moreover, the process by which people make such deferral decisions remains poorly characterized. Here, we seek to address these questions about decisions to defer, while also examining whether deferral leads to overall better decisions.

Method: Participants in two studies (lab: $N = 18$; online: $N = 38$) performed a value-based choice task in which they were allowed to defer their choices among four consumer goods. Participants viewed choices sets with varying overall and relative values (as determined by earlier valuations of individual items). They could decide whether to choose a product at that time, or whether to defer and return to that choice later. After making all of these initial choices, they then revisited all deferred choices to decide amongst them. We modeled participants' deferral decisions and decision times with a drift-diffusion model, examining how evidence accumulation toward deferral depended on the overall value and value difference of choice sets.

Results: Participants deferred on average 19% (lab) and 14% (online) of their choices. Across both studies, choices were more likely to be deferred the lower the value difference and the lower the overall value of the choice set ($OR = 0.74$; 0.39 , $ps < .001$). We found that deferral decisions were fastest when the overall value of the choice set was especially high or low, and slowest for moderately valued sets ($\beta = -0.72$, $p = .025$), accounted for by a model in which overall value determined the rate of evidence accumulation toward deferral (vs. choosing immediately). Finally, deferred choices were less consistent than non-deferred choices ($OR = 0.77$, $p < .001$).

Conclusion: Choice deferral is ubiquitous and its negative consequences are well known, but why and how people decide to defer their choices remains unclear. Using a novel task, we showed that deferral is guided largely by a dynamic evaluation of the choice set as a whole. Participants accumulate evidence about the overall value of their options and defer choosing if that value falls below a certain criterion, resulting in the fastest deferral choices when the set value is especially high or low. Examining potential explanations for such behavior, we find that deferral does not lead to better-quality choices (rather, the opposite).

P2-I-29 - A Bayesian noisy-memory account of recency effects in averaging tasks

Arthur Prat-Carrabin¹, Michael Woodford¹

¹Columbia University

Objective: Decisions must often rely on several pieces of evidence that are presented sequentially. Human choices typically exhibit "recency effects", the tendency to underweight past observations relative to more recent ones. A popular phenomenological model of recency effects consists in the exponential discounting of past observations. Why the weights of successive observations should follow this pattern, and whether they actually do, is unclear. We suggest instead a theoretical account, whereby the imprecision (noise) in the brain's recalling of past observations result in their underweighting in decisions. We investigate this hypothesis and compare its predictions to behavioral data.

Methods: We study a noisy-memory model for the estimation by a decision-maker of the average of several numbers presented sequentially. As the sequence of numbers unfolds, she maintains a running estimate of their average, but at each trial she only receives a noisy signal about her previous estimate. She however optimally decodes this signal through Bayesian inference, before updating her estimate of the average with the new observation. This model makes specific predictions as to the weights of successive observations in the final estimate, which we compare to those of an exponential-discounting model. To examine how well the two approaches fare in accounting for human behavior, we run an experiment, in which subjects ($N=63$) observe two streams of numbers and choose the one

with the higher average. We also test our model against the data of another experiment, conducted online by Clarmann von Clarenau, Appelhoff, Pachur, and Spitzer (2022), in which subjects (N=315) compare the average of one stream of number to a fixed reference value.

Results: The two experiments involve a total of six experimental conditions that differ by the task but also by the width and the shape of the prior distribution from which numbers are sampled. In all six conditions, the noisy-memory model outperforms the exponential-discount model (higher likelihood, lower BIC). The weights of recent numbers decrease more sharply in our model than with exponential discounting, while the weight of distant numbers decrease more slowly. Letting each weight be a free parameter reproduces this pattern, which our model captures parsimoniously.

Conclusions: We find, across several experimental conditions, that recency effects are better captured by a model of noisy memory than by an exponential-discounting model. Our approach illustrates how suboptimalities in human decisions can be successfully understood as resulting from the imprecision in the brain's representations of the relevant variables.

P2-I-30 - Different sources of cognitive noise explain distinct and opposing contextual modulation effects in value-based decision-making

Bo Shen¹

¹New York University

Objective: Cognitive noise is ubiquitous in decision-making. It remains unknown whether different sources of noise have different effects on value coding and choice behavior. Here we study two types of noise: noise in value representation and noise in option selection. We examine the impact of noise on context-dependent choice behavior, i.e., the impact of a third option on the choice between the other two options.

Methods: We employed behavioral testing, theoretical simulations, and neural imaging (fMRI) to study how noise impacts context-dependent behavior. In each trial of the behavioral study (N = 57), the participants chose from three Amazon goods. The third item was presented either as a real item (low representation noise) or as a cartoon image (high representation noise). Low and high time pressure was set over blocks to induce low and high selection noise. In the simulations, we assumed representation noise as adding to the options before normalization. Selection noise was assumed as adding to the product after normalization. Normalization was implemented in a divisive manner, i.e., the option value was normalized by the summed value in the choice set. In the fMRI study, we scanned 10 subjects on a similar trinary choice task (an additional 40 subjects are planned).

Results: Behavioral data shows that higher representation noise produces a positive contextual modulation: choice accuracy between two target items significantly increases with the third item's value; in contrast, higher selection noise produces a negative contextual modulation. Simulation explains the behavioral findings: increasing the third item value under representation noise enhances SNR during normalization, thus facilitating choice accuracy; whereas increasing the third item value under selection noise reduces SNR after normalization, thus impairing choice accuracy. fMRI data shows that the BOLD signal in the ventral striatum negatively codes the interaction term between the third item's value and the targets' values, suggesting a contextual inhibition regardless of representation or selection noise. Interestingly, the BOLD signal exhibits larger noise when the representation noise of the items is higher, consistent with divisive normalization.

Conclusion: We show that a contextual item can facilitate or impair choice accuracy depending on the source of cognitive noise, reconciling controversies in the literature about the direction of contextual effects. These findings support divisively normalized value coding with two types of noises as a novel computational mechanism of contextual value coding.

P2-I-31 - Context effects in multi-attribute choice

Giulia Mezzadri¹, Michael Woodford¹

¹Columbia University

Many daily decisions, such as choosing which food to eat for lunch, involve comparing options characterized by multiple attributes (Tversky, 1972). According to normative theories, the value of an option should be independent of the choice set (Allingham, 2002). However, numerous studies have shown that valuations can depend on the other options in a choice set (Tsetsos, 2012).

In this work, we further investigate the impact of context effects in multi-attribute choice using a new paradigm. Subjects had to compare two sequences of seven numbers and choose the sequence associated with the higher or lower average; this is an example of multi-attribute choice in which the true values are unambiguous, allowing study of imprecision and bias in integration and comparison without having to infer preferences. In each frame, the numbers from the two options were simultaneously presented. Sequences were generated such that numbers were drawn from each of the considered bins (see Figure 1 on the top). The numbers were then paired in three ways: in Condition 0, numbers from the same bin were paired together; in Condition 1, numbers from consecutive bins were paired together; and in Condition 2, numbers from bins that were set apart by a bin were paired together (see Figure 1 on the bottom).

Despite the fact that the numbers that were displayed were the same in all three conditions, we found that accuracy was significantly superior in Condition 0 than in Condition 1 and 2. This dis-confirms any model in which choice depends purely on the set of numbers in each sequence, independently of how they are paired in the two options, and indicates that attribute-wise comparisons must be part of the choice process. We compare the fit to our data of a variety of computational models involving attribute-wise comparisons, using cross-validation. We find that the best-fitting model is a more complex version of Tsetsos's selective integration model (Tsetsos, 2016), in which the weakest/strongest value is encoded less strongly than the others, depending on the framing of the comparison task.

P2-I-32 - Does affect bias investment decisions among professional investors?

Jennifer Eberhardt¹, Hazel Markus¹, Brian Knutson¹, Golijeh Golarai¹

¹Stanford University

Professional asset allocators manage more than \$69 trillion dollars globally on behalf of governments, universities, foundations, and companies. To perform their fiduciary duties in providing high rates of returns for their beneficiaries, asset allocators pursue various investment opportunities including for-profit financial intermediaries that are managed by professional fund managers, impacting the flow of capital across global markets. Yet, little is known about asset allocators' decision-making processes. Are asset allocators' decisions primarily driven by rational consideration, or shaped by implicit biases and affective processes that could detract from rational

choice? Does professional expertise shield asset allocators' rational choices from affective processes that typically impact decision making among lay individuals?

We designed an investment task where at each trial participants viewed information about an investment fund and rated the fund's potential for return. Funds varied along 4 distinct levels of past performance (i.e. the only reliable information on fund level). Each fund was paired with a unique identity from 1 of 4 face types (Black/White x Neutral/Smiling). Participants rated each fund's potential for return on a 4 point scale. We posited (1) higher ratings for funds paired with smiling faces compared to neutral faces, and (2) White viewers higher rating for funds paired with White faces compared to funds paired with Black faces.

To date, data from university students ($n = 10$) and professional investors ($n = 5$). As expected, funds' past performance level impacted ratings across both cohorts ($F=19$, $P<0.001$, $n=15$), and interacted with factors of race and expression ($P_s < 0.008$). However, there were no main effects of race or expression.

We also found a cohort effect (student vs. professional participants) in fund ratings, and interactions with race and expression ($P_s < 0.001$). Within each cohort, facial expression impacted ratings ($P_s < 0.03$) and interacted with race of faces ($P_s < 0.046$). Surprisingly, the effect of race of faces was evident only among the cohort of White professional investors ($F = 63$, $P = 0.004$, $n = 5$), but not among students.

These findings suggest that affect bias financial decision making across a range of professional expertise. Paradoxically however, racial bias was especially pronounced among professional asset allocators but not college students. These findings suggest diverging attitudes of college educated youth in contrast to older professionals whose attitudes may be shaped by long-term immersion in a racially homogenous professional milieu.

P2-I-33 - When models matter: How we learn to arbitrate between model-based and model-free reward learning

Leslie Held¹, Wouter Kool², Senne Braem¹

¹Ghent University, ²Washington University

Objective: In models of reinforcement learning, a common distinction is made between model-based and model-free learning, capturing how much people rely on the specific actions that led to previous rewards vs. simulations from an internal model of the environment. While there is a body of research focussing on interindividual differences in using either strategy, we aimed to test if people can learn to regulate which strategy to use based on environmental demand. Concretely, we were interested in whether exposure to lower demand (frequently repeating) or higher demand (frequently alternating) first stage states in the two-step task leads to more model-free vs. model-based control. Whereas a cost benefit analysis should favour a less demanding model-free strategy when the first states repeat and no generalization is necessary, a model-based strategy should pay off when first states alternate frequently, i.e., when generalization of information across these first states is profitable.

Methods: To this end, we adapted Kool, Gershman and Cushman's (2017) two-step task in which participants choose between four different space ships presented in pairs (stage 1) which deterministically lead to one of two planets (stage 2) with slowly drifting rewards. Importantly, in our study, participants ($n=140$) were either exposed to more (80%) first stage alternations or repetitions in an initial learning phase of 125 trials. Subsequently, they underwent a test phase of similar duration in which they were exposed to 50% first stage repetitions and alternations respectively, to which we fitted a dual systems reinforcement learning model. Our main parameter of interest was the weighting parameter w of model-based behaviour, ranging from 0 (model-free) to 1 (model-based).

Results: Our results suggest that, in line with our hypothesis, participants exposed to more first stage alternations in the training phase were more model-based ($Mw=0.79$, $SD=0.25$) in the subsequent test phase than participants rewarded more on first stage repetitions ($Mw= 0.67$, $SD= 0.32$, $t(128.97) = -2.374$, $p < 0.01$, Cohen's $d = 0.40$).

Conclusions: These findings are in line with an emerging learning perspective on cognitive control and decision making. They extend Kool et al.'s (2017) results in suggesting that the regulation of reinforcement strategies is subject to a cost-benefit analysis that, besides explicit reward cues, can also learn from experienced, environmental demands.

P2-I-34 - Individual differences in effort-based decision-making: dorsomedial prefrontal cortex metabolic and fMRI underpinnings

Nicolas Clairis¹, Arthur Barakat¹, Carmen Sandi¹

¹Behavioral Genetics Laboratory (LGC)

There are important individual differences in motivation, and these differences are key to individuals' well-being. Motivation can be defined as a series of value-based decisions which drive behavior based on a cost-benefit tradeoff between the costs and the expected benefits of an action. Despite considerable advances in the last decades in the neural correlates of value-based decision-making, the biological underpinnings that determine individual differences in effort-based decision-making remain unclear. Increasing evidence is providing strong support to implicate brain metabolism in different motivational components. However, the contribution of specific metabolites in brain regions crucially involved in decision-making to brain function and behavior in human subjects is still rather unexplored. Here, we focused on: i) the dorsomedial prefrontal cortex (dmPFC), given its key role in promoting effort-based decision-making in healthy individuals, and ii) the glutamate/glutathione (Glu/GSH) ratio. Glu is the most prominent excitatory neurotransmitter. Elevated glutamate levels can be neurotoxic, leading to enhanced levels of reactive O₂ species (ROS). On its turn, GSH has a neuroprotective role as the main antioxidant in both body and brain. Accordingly, we hypothesized that the Glu/GSH ratio in the dmPFC can explain individuals' sensitivity to exert effort by affecting dmPFC activity during decision-making. To this end, participants ($n=75$, 40 females) were first scanned using 1H-magnetic resonance spectroscopy at 7T to obtain metabolite levels in the dmPFC. Subsequently, they performed a binary choice task in fMRI, involving monetary incentives in the form of rewards and punishments and effort costs in the form of physical and mental efforts. Choice-locked neural activation in the dmPFC tracked the magnitude of the effort chosen, both in the physical and in the mental subtasks, independent of monetary stakes. Our results supported the predicted mediation model. Specifically, dmPFC BOLD responses were related to pre-task levels of dmPFC-Glu/GSH ratio, and participants' sensitivity to effort was strongly anticorrelated to dmPFC BOLD response. Accordingly, the dmPFC-Glu/GSH ratio was predictive of effort aversion, mediated by its impact on dmPFC BOLD response to effort. Our findings strongly support a key contribution of the dmPFC metabolic state to individuals' tendency to exert incentivized effort. These insights can pave the way for the development of new approaches -e.g., antioxidant treatments- to ameliorate motivational dysfunctions.

P2-K-35 - Altruism and attention to social information

Miruna Cotet¹, Mariana Blanco², Ian Krajbich^{1,3}

¹Ohio State University, ²Universidad del Rosario, ³University of California, Los Angeles

Objective: Our decisions are often influenced by what those around us do. But not everybody conforms. It is unclear whether people don't conform because they are unaware of social norms, or because they disregard them. Using eye-tracking, we examined how much people attend to and are affected by others' behavior.

Method: We investigated the effects of social influence on altruism using Dictator Games. Subjects decided how much to give while observing eight prior subjects' actions. Each subject (including the current one) occupied one of nine rectangles arranged in a 3x3 grid on the screen. The locations of the other eight subjects, as well as the mean exchange rate, remained constant within a block of 20 trials. Each block had a different mean exchange rate, as well as a new set of eight other subjects to learn about. The median giving behavior of the prior subjects varied such that different social norms prevailed across blocks.

Results: Social information affected giving rates ($b = 3.63$, $SE = 0.55$, $p < .001$), with immediate neighbors having a higher impact ($b = 3.05$, $SE = 0.58$, $p < .001$). Higher variation in social norms led to decreased giving ($b = -1.24$, $SE = 0.58$, $p = .034$). We did not find evidence of donor fatigue ($b = -0.16$, $SE = 0.27$, $p = .544$).

We characterize the type of conformity by comparing the subject's choice at the beginning of the block, prior to social information, to their later choices. Most subjects were conformists (74 conformists versus 25 anti-conformists) and were more likely to follow the group norm when it resulted in lower giving than when it required higher giving (57 selfish conformists versus 16 prosocial conformists).

We found that non-conforming subjects still looked at others. Conformity was not significantly correlated to interest in social information. Moreover, prosocial and selfish conformity were not related to gaze patterns to similar individuals.

When predicting giving rates, 30% of subjects had choices that were best fit by a model that incorporated where they looked, another 40% of subjects by a model that only included the giving of their immediate neighbors (suggesting covert attention), and the remaining 30% by a model that put equal weight on all the other subjects.

Conclusion: We used eye-tracking to investigate how gaze patterns to social information are related to altruism. We found that most subjects conform to some degree, and most conform more to a selfish norm compared to a prosocial norm. Some non-conformers are indeed aware of social norms and simply choose not to follow them. At the same time, accounting for attention does improve our ability to predict the extent of conformity.

P2-K-36 - Segmenting Experience into Generalizable Predictive Knowledge

Euan Prentis¹, Akram Bakkour

¹University of Chicago

Objective: Effective decision-making requires learning not only which choice outcomes are good, but also whether they are likely.

These likelihoods can be thought of as probabilistic event transitions that are encoded in memory as predictive representations. In the real-world, events are composed of many features. Features recur more frequently than any event they compose. Thus, predictive representations learned at the feature-level across many events are more stable, and may be re-combined in novel contexts to support generalization. The present research addresses how such feature-based representations are learned. Namely, to what extent are features segmented into independent representations (segmented learning) versus being bound in single specific representations (conjunctive learning)? We characterize how individual differences in segmentation ability relate to generalization performance.

Methods: Participants ($N = 59$) completed a task in which they made choices between multi-feature items. During training, predictive structure (start->successor) and reward value (successor->\$) could be learned at the item- or feature-level: conjunctive and segmented learning, respectively. A subsequent test then probed the ability to infer the distal reward value of start items (start->\$). Test items were either old (from training) or novel (new feature conjunctions). Generalization was instrumentalized as the accuracy on novel minus old trials. We hypothesized that when attention is highly distributed across features during training, conjunctive learners would acquire more specific representations that generalize more poorly than those acquired by segmented learners. To assess this, we fit computational models to identify conjunctive versus segmented learners, and measured attention during learning using eyetracking.

Results: Novel accuracy was above chance for segmented learners ($M = 0.593$, $SD = 0.183$, 94% HDI = [0.239, 0.933]) but not conjunctive learners ($M = 0.126$, $SD = 0.266$, 94% HDI = [-0.344, 0.657]). Moreover, in line with our hypothesis, we found a trend for the effect on generalization of the interaction between learning mechanism and attention: participants with more distributed attention generalized better when relatively better fit by the segmented model ($M = 3.846$, $SD = 2.792$, 94% HDI = [-1.276, 9.227]). **Conclusions:** These results provide initial evidence that individual differences in the ability to segment ongoing experience into features impacts generalization. Continuing work will bolster the robustness of these results, increasing our sample size (target $N = 100$) and probing generalization in additional composition and drawing tasks.

P2-K-37 - Supernormal Stimuli and the Peak-Shift Effect in Humans

Filip-Mihai Toma¹, Zhenlin Kang¹, Sarah Tashjian¹, Shinsuke Shimojo¹, Dean Mobbs¹, Colin Camerer¹

¹California Institute of Technology

Supernormal stimuli (SNS) are artificial stimuli that possess exaggerated levels of highly valued attributes found in nature, often beyond the objective value of the natural stimuli themselves. Early experiments showed that non-human animals exhibited sub-optimal preferences toward SNS that were biologically inferior to natural objects. For example, female gulls brooded giant synthetic eggs in preference to their own smaller eggs. An open question in social science is whether and to what extent humans are attracted to SNS. Given the widespread presence of SNS in various domains (e.g., social media, fast food), it is important to understand how biological preferences for SNS lead to sub-optimal choices, as SNS are preferred to similar natural objects but are inferior as choices.

Objective: Our aim is to verify if human SNS attraction can be captured through a "peak-shift" effect, which occurs when a peak frequency of choice (or subjective expected value) is shifted away from the normatively best object.

Methods: We test human SNS preference using a rational inattention (RI) model that characterizes how people allocate attention and make payoff-relevant decisions when acquiring information is costly. Imperfect information can lead to suboptimal choices that depend on the payoff of the true state and on the discriminability between "neighboring" states. Further, we run an online experiment with 60

participants to verify whether the rational inattention model can explain what humans choose when presented with perceptually close face images varying in degrees of warmth and competence.

Results: We first show that the peak-shift preference for SNS can be explained by a neighborhood-based RI model, capturing when it is more mentally costly to distinguish objects that are perceptually similar. We show that subjects choose SNS face images associated with inferior positive payoffs more frequently than the highest-payoff images, in order to avoid the negative monetary outcomes from images that are perceptually close to the highest-payoff ones. This result does not hold for the widely used Shannon entropy measure of information cost, but it does hold for “neighborhood costs” or recently developed “LLR costs”.

Conclusions: Our findings capture a peak-shift effect and provide insights into human preferences for SNS through an experiment involving perceptually similar face image stimuli. Because the concept of SNS originated in biology and it can apply to a wide range of artificially created goods, our paradigm can be extended across domains. Potential applications for understanding human SNS choice include frauds, voters’ preferences, and social media engagement.

P2-K-38 - Cognitive and neural bases of salience-driven incidental learning

Sai Sun¹, Hongbo Yu², Shuo Wang³, Rongjun Yu⁴

¹Tohoku University, ²University of California, Santa Barbara, ³Washington University in St. Louis, ⁴Hong Kong Baptist University

Objective: Humans may learn arbitrary rules by associating the current outcome with future behavior adjustments. How it occurs may depend on stable individual preferences and contextual factors, such as visual salience. Under the hypothesis that incidental learning may recruit habitual and goal-directed processes represented by modulation in attention and subjective valuation, we examined the behavioral and neural bases of task-irrelevant visual salience-guided outcome evaluation and behavioral adjustments.

Method and Results: We employed a simple binary choice gambling task and quantified the frequency at which participants switched to an alternative option after observing the feedback (Exp.1-6, total N=184). We first established the baseline behavior without salience in Exp.1 (n=21). We then highlighted the utility or performance dimension of chosen outcome using colors in Exp.2 (n=30) and demonstrated that the switching difference increased along the salient dimension, confirming a salience effect. Moreover, the salience effect was abolished when directional feedback information was removed in Exp.3 (n=28), suggesting that the salience effect is feedback specific. Using text emphasis combined with eye-tracking, we further generalized our findings and examined the attentional basis of feedback-specific salience effects. The fixation differences between chosen and unchosen values were enhanced along the feedback-specific salient dimension (Exp.4, n=48) but unchanged after removing feedback-specific directional information (Exp.5, n=32). Moreover, salient dimension-driven behavioral switching was correlated with fixation properties, suggesting that feedback-specific salience guides attention deployment and further drives behavioral adjustments. Lastly, our fMRI study (Exp.6, n=25) showed that the striatum subregions encoded salience-based outcome evaluation, while the vmPFC encoded salience-based behavioral adjustments. The connectivity of the vmPFC-ventral striatum accounted for individual differences in utility-driven, whereas the vmPFC-dmPFC for performance-driven behavioral adjustments.

Conclusions: Our findings delineated how specific salience modulates feedback-based attention deployment and further guides behavioral adjustments. Our neuroimaging study further revealed the neural substrates underlying specific salience-driven learning processes involving the frontal-striatal circuit. Our results provide a neurocognitive account of how specific visual salience drives incidental learning by integrating habitual and goal-directed systems represented by attention and valuation modulation.

P2-K-39 - Brain network signature of attention during cue-approach training predicts change in food preferences

Xinyue Li, Ziwei Zhang¹, Monica Rosenberg¹, Akram Bakkour

¹University of Chicago

Objective: Targeting automated memory and attentional processes to achieve behavioral change may be an effective approach to address public health concerns such as obesity. Previous research has shown that a simple task (cue-approach task or CAT), which does not rely on external reinforcements, can shift choice behavior by targeting specific items for approach during training. Studies using CAT have highlighted the importance of sustained attention to particular items to induce a change in choice behavior. However, the task design did not allow for a precise behavioral measure of sustained attention during training. In the current study, we aimed to explore the role of sustained attention during CAT using a network-based predictive model of attention applied to neuroimaging data.

Methods: We re-analyzed CAT fMRI data with food items (N = 31) using an edge-centric brain network method. This validated method identifies timepoint-to-timepoint co-fluctuations between all pairs of nodes in a 268-node brain atlas by computing the element-wise product of their normalized BOLD time series. This resulted in a 268x268 co-fluctuation matrix at every time point. We applied the pre-trained sustained attention connectome-based predictive model (saCPM) to the edge-centric co-fluctuations to get a time-resolved index of sustained attention. The saCPM consists of 2 large-scale functional networks, the high- and low- attentional networks, whose strength predicts better and worse sustained attention, respectively. Mean co-fluctuation in high- minus low-attention networks was our measure of predicted sustained attention.

Results: Contrary to our prediction, we found no significant difference in average predicted attention scores for Go items (those targeted for approach) and NoGo items, though the score for Go was numerically higher than for NoGo training trials (mean difference = 0.003, t = 1.29, p = 0.20). We also found no relationship between average predicted attention during training and participants’ later choice behavior ($\beta = 5.68 [-5.12 - 16.47]$, p = 0.30). However, our result showed that greater item-level increase in predicted attention as training progressed was related to the shift in preferences for Go items ($\beta = 4.49 [0.11 - 8.85]$, p = 0.04).

Conclusion: The degree of sustained attention to Go items predicts the shift in those items’ subsequent appeal. The results highlight the importance of targeted attention in shifting item preference during a training regimen that does not rely on external reinforcements.

P2-K-40 - Differential effects of prior outcomes and pauses on the speed and quality of risky choices

Zhang Chen¹, Charlotte Eben¹, Frederick Verbruggen¹

¹Ghent University

Failures to obtain rewards can influence both the direction (choosing what to do) and intensity (response vigor) of subsequent motivated behavior. For instance, in risky decision-making, losses can induce faster responses (‘intensity’) and sometimes increase risk-taking (‘direction’), which may lead to detrimental consequences in some situations. Inserting pauses have been proposed to

counteract these motivational influences in some responsible gambling tools. However, the joint effects of prior outcomes and pauses on these two components in risky decisions have been unclear. To examine this question, we conducted four online experiments ($N = 47, 58, 127$ and 181 , respectively), in which participants (from the general population on Prolific.co) alternated between a guess game, in which they won or lost money via guessing, and a choice game, in which they chose between two risky options with different expected values. Across trials, the pause between a guess game and a choice game was made either short (0 millisecond in Experiments 1 and 2, or 300 milliseconds in Experiments 3 and 4) or long (3000 milliseconds). In all four experiments, prior outcomes consistently influenced decision speed, so that participants chose more quickly after a loss than after a win. However, prior outcomes did not systematically influence people's choices. In contrast, pauses increased decision quality in risky choices, so that participants chose the option with a higher expected value more often, while their decision speed was not substantially impacted by pauses. Pauses may improve decision quality by increasing attention to relevant information, as its effect was absent when the overall task attention was increased (Experiment 3). These findings have important implications for both responsible gambling that makes use of pauses in play, and risky decision-making research that has neglected the potential influence of inter-trial intervals so far. Future work can examine the underlying computational and cognitive processes, and the generalizability of these findings to other contexts and populations.

P2-L-41 - Neurocomputational mechanisms of sensory plasticity associated with the one-armed bandit task

Anush Ghambaryan¹, Anastasia Grigoreva², Anna Shestakova², Victoria Moiseeva³

¹HSE University & École normale supérieure, ²HSE University, ³Higher School of Economics

Several studies have shown plastic changes in the human sensory cortex associated with accumulation of experience, particularly, with repetitive behavior (Shuler & Bear, 2006; Vickery et al., 2011). Such neuroplasticity implies monetary valuation-driven alterations in the stimuli perception (Persichetti et al., 2015; Thomas et al., 2013), where valuation of stimuli with regard to monetary rewards and losses is designed in the monetary incentive delay task (MID) paradigm (Knutson et al., 2000; Lutz & Widmer, 2014; Oldham et al., 2018). The latter however does not meet conditions and assumptions of computational algorithms of Markov decision processes (MDP) widely used for accessing learning and valuation brain networks (Daw et al., 2005; Gläscher et al., 2009). Hence little is known about computational mechanisms of neuroplastic changes along with value learning.

Based on evidence that learning in cortico-striatal synapses projecting information from the auditory cortex to the auditory striatum is driven by reward contingencies (Elias et al., 2015; Xiong et al., 2015), Gorin et al. (2020) hypothesized and showed that value-based decisions modulate perceptual processing. In this study, mismatch negativity (MMN) signals of the auditory cortex (roving oddball paradigm) were tested for their association with monetary outcomes (MID paradigm). In our study, we have integrated the MMN paradigm with the one-armed bandit paradigm (Sutton & Barto, 1998). Specifically, we have designed a task where subjects are exposed to pairs of auditory stimuli for generating MMN signals while they engage in MDPs driven by differential monetary values of the auditory stimuli.

We probed the newly designed task in 37 subjects and found that the performance measured in the choice rate of the option with the highest monetary outcome was above the chance level ($T(36) = 12.7, p < 0.00001$). Next, we modeled the behavioral data by applying the Rescorla-Wagner learning rule (Rescorla & Wagner, 1972) separately for two auditory stimuli carrying low and high monetary outcomes. According to preliminary results of model fitting, there is significant difference between the two learning rates across subjects ($T(36) = 4.5, p = 0.0001$), which is in support of the developed paradigm for our study purposes. Further in the neural analysis we will associate the learning rates and parameters of soft-max choices with the MMN signals. This will allow us to describe and analyze neuroplastic changes in the auditory cortex as a reinforcement learning process.

P2-L-42 - Neuroeconomic data for the calibration of disequilibrating excessive demand in economics and finance models

John Haracz¹, Mehreen Zakaria²

¹Indiana University, ²Columbia University

Objective: Dynamic stochastic general equilibrium models have been widely criticized for failing to forecast the Global Financial Crisis of 2008-09 (Yellen, 2010; Guzman & Stiglitz, 2020; Vines & Wills, 2020). This and other flaws of neoclassical economics were proposed to arise partly from the failure of equilibrium-based models to capture excessive demand (Haracz, 2021a,b, 2022a,b), which exceeds the balanced "excess demand" in general equilibrium theory (Arrow, 1974, p. 266; Debreu, 1984, p. 270). Excessive demand is defined as demand that promotes disequilibria in asset or goods markets and drives prices above fundamental values. The present theoretical study seeks: 1) potential neuroeconomic biomarkers of excessive demands; 2) to develop modeling methodology that calibrates excessive demand from neuroeconomic data.

Methods: A systematic literature review focused on neuroeconomic studies of tasks that elicit demands (i.e., choices) by subject groups that may be prone to show excessive demands (e.g., individuals with substance use disorders, eating disorders, or gambling disorder, as well as investors or lab-market subjects who trade excessively). The studies' results were assessed for their potential to enable the calibration of excessive demand in an adaptation of the dynamic disequilibrium model with randomness (Guzman & Stiglitz, 2020; Stiglitz & Guzman, 2021; henceforth "GS model") for analytic and policymaking applications in macroeconomics and finance.

Results: Compared to healthy control groups, the above subject groups typically show altered fMRI-measured activity in the brain reward system (e.g., ventral tegmental area, nucleus accumbens, and ventromedial prefrontal cortex) during anticipation, choice, or outcome phases of demand-elicitation tasks. These activity changes (e.g., Smith et al., 2014; Moeller et al., 2018) show a potential for calibrating excessive demand in the GS model, especially when accounting for fMRI measurement error in predicting demand (Webb et al., 2021). To model excessive demand, the GS model's t variable, which denotes downward pressure on demand (Stiglitz & Guzman, 2021, p. 471), could be supplemented with another L variable, denoting upward pressure on demand.

Conclusions: Neuroimaging research may yield biomarkers of traders' or consumers' excessive demands. A high biomarker prevalence could indicate that financial- or commodity-market demands have exceeded boundary conditions, beyond which equilibrium-oriented models are less applicable than alternatives (e.g., novel disequilibrium [Guzman & Stiglitz, 2020], multiple equilibrium [Vines & Wills, 2020], or temporary equilibrium models [Brunnermeier et al., 2021]).

P2-L-43 - The interplay of hedonic appetite and attentional abilities is linked to poorer dietary self-control

Łukasz Tanajewski¹, Todd Hare², Jakub Skalbani^{1,3}

¹Kozminski University, ²University of Zürich, ³University of Warsaw

Aims: Voluntary attention supports dietary self-control, while hedonic appetite impairs it. Neuroimaging and behavioral experiments have suggested that hedonic processing of foods may consume voluntary attention. We therefore checked whether and how the interaction between hedonic appetite and attentional abilities is associated with poorer dietary self-control.

Methods: Using the Power of Food Scale (PFS) we measured hedonic appetite in two samples of young adults. In Study 1 (380 participants) poor dietary self-control and attentional abilities were measured by the Uncontrolled Eating subscale of Three-Factor Eating Questionnaire and the Attentional Control Scale, respectively. In Study 2, 49 participants rated the healthiness and tastiness of foods. In the incentivized series of computerized trials, dietary self-control was defined as (i) choosing a healthier over a tastier food item (a binary measure), and as (ii) the probability of dietary self-control, i.e., the probability that a healthier and less tasty food item was chosen (a continuous measure). Measure (ii) accounted for the imprecision of the tastiness and healthiness ratings. Attentional abilities were computed as the ratio of mean to standard deviation of response time in a memory task check (a presented digit was checked against a memorized digit). Multiple linear ordinary least squares regressions were applied in Study 1. Multilevel mixed-effect regressions were used in Study 2. Simple slopes analyses were employed to explore how the interaction of hedonic appetite and attentional abilities was associated with dietary self-control.

Results The interaction between the PFS and attentional abilities was associated with poorer dietary self-control (Study 1: standardized beta (SB)=-.54, $p=.02$; Study 2: OR=.64 ($p=.02$) (measure i); SB=-.71 ($p<.01$) (measure ii)). The strength of the positive association between attentional abilities and self-control decreased with a PFS score, with no association (Study 1) and the negative association (Study 2) for high-PFS individuals. The strength of the negative association between the PFS and dietary self-control increased with attentional abilities, with no association for individuals with low attentional abilities in Study 2.

Conclusions: The interplay between attentional skills and hedonic appetite can be harmful for dietary self-control in young adults. A single-self perspective to study self-control developed in neuroeconomics, rather than dual-self and ego depletion perspectives, can be helpful in explaining our findings. If hedonic appetite is strong enough, attentional abilities may impair dietary self-control skills.

Funding: The study was supported by Polish National Science Center Grant 'Cognitive resources and self-control of impulses: a neuroeconomic approach' (No. 2017/26/D/HS6/01157). PI: Łukasz Tanajewski.

P2-L-44 - Disentangling the affective impact of stimulus proximity and motion using immersive virtual reality

Tara Srirangarajan¹, Yan Yan¹, Jeremy Bailenson¹, Brian Knutson¹

¹Stanford University

Anticipatory affect, or the emotional states that people experience while anticipating significant outcomes, guides motivated behavior and predicts choice. In recent years, Virtual Reality (VR) technology has been leveraged as a tool for mapping emotions to physiological measures, such as heart rate and pupillometry. However, prior studies have predominantly measured affective states only after the fact. In this study, we aimed to induce and track affect on a second-to-second timescale (coupled with real-time psychophysiological measures), with the goal of elucidating how basic visual characteristics influence affective responses in immersive VR. We designed a variant of the Monetary Incentive Delay (MID) task, which has been shown to reliably elicit changes in self-reported affect and neural activation in response to incentive cues. Subjects ($n=30$) completed a virtual adaptation of the MID task, which was projected using an HP Reverb G2 Omnicept headset. The task consisted of 96 trials and manipulated three within-subject variables: cue magnitude (six conditions: $-\$5$, $-\$1$, $-\$0$, $+\$0$, $+\$1$ and $+\$5$; cue type (four conditions: proximal, distal, looming and receding); and affective probe (two conditions: anticipation probe or outcome probe). On each trial, participants first saw an incentive cue (2s), followed by an anticipation period (~2s). A target then appeared (350-550 ms), and participants attempted to "hit" it by pressing a button on the hand-held controller. This was followed by an outcome period (2s) and a variable inter-trial interval (2, 4, or 6s). Affect probes were inserted in either the anticipation or outcome phase of each trial, during which participants reported their momentary affective valence and arousal. Behavioral results suggested that incentive cues influenced affective arousal and valence, such that distal cues had less of an affective impact compared to all other cue variants, and that this effect was not as influenced by cue motion. This finding was also consistent with increased pupillary dilation as a function of incentive cue magnitude, specifically during the anticipation phase of each trial. For the outcome phase, we found that incentive outcomes drove affective valence, such that hits increased, while misses decreased valence. This effect was not influenced by cue proximity or motion. Overall, these results suggest that cue proximity but not motion modulates anticipatory affect. Thus, these multi-modal findings represent a key methodological and conceptual advance towards translating neuroeconomic research into new mediums including self-report and physiology to better understand dynamics of affect and motivated behavior.

Poster Session 3

P3-C-1 - Inferring latent social network properties aids strategic gossip-sharing

Alice Xia¹, Apoorva Bhandari¹, Oriel FeldmanHall¹

¹Brown University

Knowing who is friends with who in a social network can inform how we make decisions about others, such as determining which friends to trust with gossip (i.e., discussion about third parties who are not present). Remarkably, gossip often stops short of reaching the person it is about (Blumberg, 1972), which suggests that people are strategic in how they share sensitive information. What types of social network information do people rely on to optimize strategic gossip-sharing? How might multidimensional information help or hinder this process? Here we assess which social network properties people leverage to build mental representations that enable strategic decision-making. In three experiments, participants were tasked with achieving the common phenomenon of sharing gossip widely within a network without it reaching the target of gossip. Participants learned about an artificial social network and then rated how likely they would be to share gossip with each network member, given a target of gossip, an individual who was also embedded in that network. In Exp. 1 (N=202), participants had to build mental representations of the social network from friendship information alone. To test how people's network representations might change in the presence of additional social information, in Exp. 2 (N=199), participants learned about the trustworthiness of each network member prior to learning about friendships. Finally, to more closely reflect how social information is encountered in the real world, participants in Exp. 3 (N=176) were exposed to multidimensional information and were required to simultaneously learn about friendships and each individual's trustworthiness. Across all experiments, we found that people: (1) encode friendship information about individuals in a social network and subsequently use it to avoid sharing gossip with the wrong person (i.e., friends of the target of gossip); (2) infer goal-relevant, but latent topological properties of the network – such as popularity and the distance between two network members – to strategically gossip more with popular individuals far away from the target, while avoiding the target's friends or friends-of-friends; (3) use an individual's trustworthiness to fine-tune who to gossip with. While people can gossip strategically by building mental representations of a densely connected social network based on piecemeal observations, learning multidimensional information degraded the fidelity of these mental representations. Overall, our results show that latent network properties, combined with individual-level trait information, guide goal-oriented choice in complex social environments.

P3-C-2 - Diversity Matters: An Agent-Based Simulation Approach on Group Cooperation and Performance

Chris Stolle¹, Bartosz Gula², Yi Huang¹

¹Lingnan University, ²University of Klagenfurt

Objective: Globalisation and technological possibilities allow people to collaborate across large distances, and work processes are developing from homogeneous individual tasks to more complex group achievements. These developments make people from different backgrounds and with different abilities work together. In short, diversity increases.

Many studies have theoretically and empirically shown that diversity positively influences team outcomes. However, little attention has been paid to the diversity effect on the functionality of teams. This study used an evolutionary simulation approach to investigate diversity effects on group cooperation and performance.

Methods: We identified three kinds of diversity commonly reported to improve group performance: (1) member variability of accuracy in doing a specific task, (2) individual decision-making strategies, implemented as the variety of heuristics used, and (3) information sources as access to different but equally good information. Each of these types of diversity was formalised in a group foraging task setting: A group selects among several places to forage and distributes the food found equally among all members. The places contain different amounts of food, that can be estimated from a noisy cue set. The simulated agents have the following task: First, they decide whether to pay cooperation costs to observe the cues or save the costs and rely on others' decisions (i.e. free ride). Next, cooperating members choose a place by majority voting.

For cooperation, not only the social aspect (i.e. the group) plays a role, but also the (simulated) environment. We assessed diversity effects across several critical aspects of the environment: (i) the number of choices that are available for foraging, (ii) cooperation costs, (iii) redundancy of the cue structure and (iv) the compensatory level of the cues.

Results: The results show that across all environmental conditions diversity in decision-making and information sources increase cooperation rates. Diversity in accuracy did not affect the cooperation rate. Group performance depends more on a diversity X environment interaction, with more positive diversity effects on performance in non-compensatory than in compensatory environments.

Conclusion: The simulation confirms the general observation of positive diversity effects on teams. Further, we dissociated the effects of specific kinds of diversity on team performance and cooperation rates. Therefore, our simulation results can be used as a framework to classify diversity effects from empirical studies.

P3-C-3 - Aligning of minds: Shared neural codes predict online coordination

Cong Wang¹, Lusha Zhu¹

¹Peking University

Objective: Coordination—selecting complementary decisions for mutual gains—is an important form of cooperation. A large body of research in game theory and multi-agent AI proposes that coordination may involve complex, recursive reasoning about the mental states of social partners. In humans, such decision processes are often cognitively demanding and error-prone, raising questions regarding how humans achieve effective, flexible everyday coordination in a seemingly effortless manner. Here we explore neural substrates for a computational shortcut to coordination, building on a long-standing hypothesis that coordination may be supported by the common knowledge of a focal point that people tend to choose by default. Combining the classic Pure Coordination Game with inter-subject analysis of fMRI pattern correlation, we tested whether the neural alignment of knowledge representation could predict coordination behavior, reliably and specifically, across tasks and subjects—a necessary condition for the theory of common knowledge.

Methods: Thirty-four subjects underwent fMRI scanning while they were viewing a sequence of landscape/cityscape images passively, without making any decisions. Images used in the fMRI experiment were randomly arranged into groups of 4 and used for an online Pure Coordination game (N = 542), where a subject needs to select one picture out of 4 to match the choices by other online subjects who are playing the same game and facing the same requirement of choice alignment.

Results: We quantified the between-subject similarity in the neural response pattern to each image and used it to predict how likely a

particular image would be selected in the coordination game. This identified a network of brain regions including the posterior cingulate cortex (PCC), parahippocampal cortex, and ventromedial prefrontal cortex (permutation test, voxel-wise FDR corrected $q < 0.05$). Importantly, the across-task, across-subject prediction was only seen in coordination, but not in a matching condition where subjects were facing the exactly same set of pictures but asked to make personal decisions without the need of matching with others.

Conclusions: The data support the account that successful coordination may involve aligning the internal representations about the external world. It points to the neural substrates associated with coding common knowledge, which likely underlies tactic agreement, convention, and norms, critical for mutualistic behavior.

P3-C-4 - Neural and behavioral dynamics of reciprocal sibling interaction in the Taylor Aggression Paradigm

Julia Koch¹, Lucia Hernandez Pena, Rik Sijben, Edward S. Brodtkin, Ute Habel², Lisa Wagels²

¹University Hospital & RWTH Aachen University, ²Medical Faculty, Uniklinik RWTH Aachen

Objective: Social interaction is a dynamic process, and its authenticity subconsciously affects our action patterns [1]. Especially siblings often show specific responses to each other due to their familiarity and type of relationship. Nonetheless, interactions have been mostly studied using fictional opponents. We here aim to study neural and behavioral underpinnings of sender-receiver exchange in a competitive context. Therefore, we introduce an fMRI hyperscanning approach capturing joint brain activity during an interactive Taylor Aggression Paradigm (iTAP) in young adult siblings. Initially used to study laboratory aggression [2,3,4], our iTAP provides the opportunity to study more complex behavioral and neural exchange dynamics. This increases the ecological validity of the paradigm and enables the detection of behavioral and neural within-dyad dynamics and between-dyad patterns. In addition, within-person and within-dyad factors influencing the exchange pattern can be examined.

Methods: Data of a behavioral study ($n = 28$ sibling pairs) and the follow-up fMRI study ($n = 46$ sibling pairs) will be included. In both studies, siblings performed the iTAP by setting an individual monetary punishment level per round and subsequently competing in a rock-paper-scissors game against each other in real-time. The money selected by the winner was withdrawn from the loser of the game round.

Groups of similar dyad-level time series were created by hierarchical clustering [5]. A repeated-measure version of the actor-partner interdependence model was employed to integrate the effect of one's previous choice as well as on the other sibling's choice for the following trial [6]. Further, within-dyad roles related to competition levels (high vs. low) were ascribed to investigate if differences in those effects would become evident.

Results: Siblings can be characterized by differences in overall monetary withdrawal, as well as mutual adaption in money level. Further, siblings exert a reciprocal influence on each other's choice behavior in the iTAP ($b = .35$, $p < .001$) and show consistency in their own behavior over time ($b = .55$, $p < .001$). Those siblings that see the relationship as more competitive are influenced more strongly by their siblings' behavior ($b = .31$, $p < .001$). Functional imaging results modelling the interactive choices will be presented to complement the behavioral model.

Conclusion: The iTAP is a valid tool to study interaction dynamics and patterns in siblings. We could observe reciprocal influences in behavioral choices and direction modulated by levels of competitive attitudes dividing siblings within dyads.

P3-C-5 - Stereotypic expectations entrench unequal conventions across generations in deep multi-agent reinforcement learning

Rebekah Gelpi¹, Yikai Tang¹, Ethan Jackson, William Cunningham¹

¹University of Toronto

Successful social coordination in large groups requires making predictions about unfamiliar others that can take the form of generalized beliefs about social groups. However, if these predictions lead others to change their behaviour in order to match the expectation, this can result in stereotypic beliefs being confirmed and entrenched. In a set of 3 studies (population size range: 30-600, 30 simulations each), we investigate the emergence and persistence of this potential "feedback loop" using a multi-agent deep reinforcement learning paradigm, in which agents must successfully coordinate with each other to achieve rewards. Crucially, agents have differing skills, so if its actions are correctly predicted, an agent will maximize its reward by taking actions that match its skillset.

We design an environment in which agents have observable perceptual cues that are correlated with their respective skillsets. Because the majority of the members have a group have a similar skillset, the group becomes associated with the respective actions. For the minority in a group who have the less common skillset, successful coordination can be achieved by either successfully predicting that some individual agents take different actions, or through these agents changing their behaviors to match the prediction.

We find that in these environments, stereotypic (cue-based) predictions result in minority agents changing their behavior to match this stereotype, even though this requires them to take actions they are unskilled at. For example, relative to an environment where agents must be learned as individuals because no correlated cue is available, agents whose skillset is more uncommon within the group engage in a lower proportion of skilled actions ($z = -102.59$, $p < .001$, Cohen's $d = -1.18$) and obtain a lower mean reward ($z = -78.70$, $p < .001$, Cohen's $d = -1.15$).

Furthermore, because this behavior confirms the prediction that agents who match a cue mostly behave in a stereotypic fashion, these stereotypic expectations carry over towards unobserved agents who match this cue. Thus, when new agents are progressively introduced to the environment with no underlying difference in skillsets – resulting in an environment in which the original cue is no longer correlated with an agent's skill – these stereotypes persist across generations of agents, resulting in a lower collective reward across all agents ($z = -5.98$, $p < .001$, Cohen's $d = -0.25$). Our results suggest that even in the absence of motivated cognition and cognitive limitations, the need for social coordination may play a role in developing and maintaining social stereotypes.

P3-D-6 - Bidding with Effort: Using Grip Force to Assess the Value of Information

Haixue Fan¹, Bryan Dong¹, Dorothy Benkelman¹, Alexandra Rodman², Hayley Dorfman¹, Paul Glimcher³, Elizabeth Phelps¹

¹Harvard University, ²Northeastern University, ³New York University

Uncertainty, prevalent in everyday scenarios, can be decomposed into risk (known probability) and ambiguity (unknown probability). Past work in decision-making has shown that people tend to avoid decisions associated with ambiguity more than risk, but questions remain about how people value the resolution of uncertainty when no direct outcome is at stake. The existing literature on information-seeking behavior usually uses people's willingness to pay to quantify the value of information. In the current study, we examine people's

valuation of information under risk and ambiguity using a novel grip force measurement that quantifies a participant's willingness to exert effort in exchange for information. In the experiment, participants (N = 12, data collection ongoing) view either a risky or ambiguous lottery, and then report their maximum willingness to exert effort to reveal the outcome of the lottery by briefly squeezing a hand dynamometer ('effort bid'). After placing the effort bid, a standard Becker-DeGroot-Marschak auction procedure is used to ensure that the subjective valuation has been elicited in an incentive compatible manner. Specifically, the participant's bid is compared to a random number ('effort price') lying between 0 and their maximum possible grip force as measured in a prior calibration session. If the bid is lower than the price, the participant loses the bid and will not see the outcome of the lottery. If the bid is higher than the price, the participant wins the bid and will pay the effort price for a consecutive 3 seconds, after which the outcome of the ambiguous or risky lottery is revealed. Each participant completed 81 trials. Our preliminary results suggest that people are willing to exert effort in exchange for information that does not impact their payoff ($\beta = 0.58$, $t(11) = 6.68$, $SE = 0.087$, $p < .001$). In the context of risk, participants are more willing to exert effort (i.e., place a higher effort bid) when the expected value of the lottery is high ($\beta = 0.02$, $t(634) = 3.55$, $SE = 0.006$, $p < .001$). In contrast, people's effort bid is indifferent to variations in ambiguity ($\beta = 0.06$, $t(311) = 0.64$, $SE = 0.09$, $p = .525$). Finally, the amount of effort people exert when they 'win' the auction positively correlates with their effort bid ($\beta = 0.62$, $t(482) = 5.17$, $SE = 0.112$, $p < .001$; controlling for the effort price), suggesting a match between explicit and implicit valuation. Together, our results reveal a dissociation between people's valuation of information resolving risk versus ambiguity in the effort domain and contributing to the growing literature in understanding the valuation process of information.

P3-D-7 - Role of Phasic Mesocorticolimbic Dopamine Release in Risk/Reward Decision Making

Jackson Schumacher¹, Mudi Zhao¹, Debra Bercovici¹, Melody Hall¹, Stan Floresco¹

¹University of British Columbia

The ability to gauge the profitability of multiple options under conditions of risk and uncertainty is critical for human economic decision making and convergent evidence points towards a key role for the mesocorticolimbic dopamine (DA) system in mediating this process. To study the underlying neurobiological mechanisms involved, risk/reward decision making is often operationalized using rodent operant tasks such as probabilistic discounting which requires rats to choose between a small, certain and a large, risky reward whose probability varies systematically across blocks of trials. Prior work using pharmacological DA receptor blockade has implicated dopamine transmission in both the medial prefrontal cortex (mPFC) and nucleus accumbens (NAc) in probabilistic discounting. However, DA release manifests as discrete phasic events and pharmacology lacks the temporal resolution to disentangle the effects of DA release on the sub second timescales on which decisions are made. Therefore, the objective of the current study was to interrogate the temporal dynamics of DA modulation in probabilistic discounting. To do so, we expressed the inhibitory opsin Arch3.0 in ventral tegmental DA (VTA) DA neurons of male and female rats to induce optogenetic silencing of either DA cell bodies or terminals in the NAc or mPFC. Temporally discrete suppression of DA neurons/terminal activity was induced during different task events, including prior to choice, and during different choice outcomes (large/risky rewards, small/certain rewards or reward omission). Preliminary data reveal that silencing VTA DA cell bodies (n=10) during receipt of large/risky reward shifts choice towards the small/certain option and, conversely, silencing VTA cell bodies during small/certain reward delivery shifts choice towards the large/risky option. Surprisingly, neither manipulation altered sensitivity to previous wins (win-stay behavior) but the two manipulations did alter loss sensitivity with large/risky reward inhibition increasing and small/certain inhibition decreasing loss-shift ratios. Silencing DA in the NAc (n=7) (but not mPFC; n = 5) after rewarded large/risky choices reduced risky choice when the odds of reward increased across a session, whereas mPFC or NAc terminal inhibition during small/certain reward delivery was without effect. No manipulation had any effect when silencing occurred during reward omissions. Together, these findings demonstrate that dopaminergic modulation of decision-making behavior critically varies along temporal and spatial dimensions such that silencing DA has differential effects in different brain regions and different behavioral epochs.

P3-D-8 - Computational Contributions to Prosocial and Utilitarian Risk-taking

Jiajin Zhang¹, Sarah Tashjian¹, Alma Merscher², Dean Mobbs¹

¹California Institute of Technology, ²University of Würzburg

Real world risky decision making often involves a social trade-off in which the decision will bring positive consequence to one agent but negative consequence to another. Previous studies have investigated factors that influence economic decisions-making, but rarely with a direct social trade-off. Further, little is known about what drives prosocial decision-making, making a choice to benefit another at a personal expense, despite the vast social consequences of prosocial behavior. The current preregistered study (<https://osf.io/ny4as/>) used a novel variant of the Columbia Card Task in which participants' decisions had opposing consequences for themselves and a charity of their choosing. Participants (N=22, ages 18-40 years) played two runs of the Columbia Card Task while wearing wrist devices that measured electrodermal activity (EDA). In one run, the score participant gained conferred direct monetary benefit to themselves, but the losses incurred were taken from a donation to charity. In the other round, the contingencies flipped such that the charity received an actual monetary benefit for gains received whereas the participant's bonus was reduced by the loss incurred. Analyses were performed in R using lme4 and general linear mixed model. We found that decision-making is predicted by context ($B = -0.351$, $p < 0.001$), risk ($B = -0.031$, $p < 0.001$), and return ($B = 0.012$, $p < 0.001$). Participants were more likely to make decisions that benefited themselves at the cost of charity; higher return and lower risk were associated with more likelihood of taking risk. Individual differences in impulsive risk-taking (BAS Fun seeking), loss aversion, and trait anxiety influenced risk-taking in the utilitarian context (gain for self, loss for charity) but had no significant influence in the prosocial context (Fun Seeking: [$B = 0.279$, $p < 0.01$]; loss aversion: [$B = -0.129$, $p < 0.05$]; Anxiety: [$B = -0.149$, $p < 0.05$]). Males were also more likely to take trade-offs at the cost of others. Interestingly, the expected outcome (i.e., return, which was operationalized as the expected value of a given decision) increased prosocial risk-taking. Further, arousal was higher as measured by EDA when subjects took prosocial risks at the cost to themselves. Our results implicate prosocial behavior may be driven by increased magnitude of the benefit to others rather than individual differences of the actor whereas self-oriented utilitarian choices are shifted by individual differences. Affective salience in physiological arousal may be one mechanism by which expected value amplifies prosociality.

P3-D-9 - Volatile stocks cause higher inter-subject correlation in the right anterior insular cortex

Kirill Efimov¹, Ksenia Panidi², Victoria Moiseeva², Anna Shestakova¹, Iiro Jääskeläinen¹, Vasiliy Kliucharev³

¹HSE University, ²Higher School of Economics, ³Higher School of Economics - National Research University

Objective: This study aims to investigate the influence of a stock price volatility on perception and behavior of the market participants. Previously, increased activation of the right lateral prefrontal cortex (LPFC) and right posterior parietal cortex (PPC) during trading has been observed (Ragetti et al., 2017), while activation of the anterior insular cortex (AIC) has been traditionally considered to play an important role in loss anticipation, risk aversion (Kuhnen and Knutson, 2005) and irreducible risk estimation (Bosschaerts, 2010). However, the effect of volatility, to our knowledge, has not been isolated.

Methods: 42 healthy participants were engaged in the imitated stock trading environment, where they observed an updating stock price graph for 60 price points followed by a series of 8 trading decisions interleaved with periods of 5 price points. All stock price graphs were constructed from publicly available histories of day-to-day stock prices of NASDAQ-indexed companies: Ten 100-day periods that exhibited random properties by the Hurst coefficient (Hurst, 1951) were extracted, centered to the mean value equivalent of 1.5\$ and varied to form 2 distinct conditions with a 10-fold difference in volatility. Participants underwent fMRI scanning sessions with the use of Philips Ingenia 3T scanner (Netherlands) while engaged in stock market imitation. The preprocessing of the data was conducted with the application of fMRIPrep (Esteban et al., 2019), SPM12 (Penny et al., 2007), and ICA-AROMA (Prum et al., 2015). An inter-subject correlation analysis (ISC) (Hasson et al., 2004; Kauppi et al., 2010) was conducted in voxel-wise manner with the use of ISC toolbox (Kauppi et al., 2014) and in ROI-wise manner (Chen et al., 2016) with the parcellation by the Brainnetome Atlas (Fan et al., 2016).

Results: The participants' observation of the graphs with higher volatility caused significantly higher inter-subject correlation in the right anterior insular cortex, right lateral prefrontal cortex and areas of parietal lobe in comparison with the low volatility condition ($FDR < 0.05$). Participants also traded more actively in high volatility condition: they bought more total shares per session on average (Wilcoxon paired test $p < 10^{-6}$) and had higher trading volume (Wilcoxon paired test $p < 10^{-4}$).

Conclusions: Our results suggest that higher volatility results in increased risk-taking behavior. Moreover, the volatility-induced processing of risk and uncertainty within AIC and PFC is synchronized between individuals.

P3-E-11 - Rational decision making for distant goals

Shiyi Liang¹, Evan Russek², Robb Rutledge³, Elliott Wimmer¹

¹University College London, ²Princeton University, ³Yale University

The process of evaluating and pursuing goals plays a significant role in important life decisions. Goal pursuit usually involves multiple steps, where, at each point, one faces some risk of failure. Accurately estimating this risk is crucial for evaluating potential actions and their outcomes. While extensive research has investigated one-step decisions, risks in the real world typically compound. In the present study, we evaluate multi-step goal-directed behavior, in an environment where risks compound and must be learned. Participants completed a learning and decision-making study across two days. We designed a novel task that required participants to utilize their memory to pursue goals by navigating a semi-circular maze track divided into 16 discrete locations. Each location was marked by an image of a unique object. After initial free exploration, participants engaged in goal-pursuit trials where they were provided with a start location and a goal location, and received a reward upon successful navigation to the goal. Participants learned only via experience that the path could be unexpectedly terminated, according to an independent risk of failure (10%) at each maze location. To assess risk-taking, participants then faced choices between a constant certain option and varying distant rewards made uncertain by the risk of failure along the way (subsequently receiving the certain reward or engaging in goal pursuit). The multi-day design allowed participants to build consolidated cognitive maps of the maze and representations of learned risk, and thus our analyses focus on day two choices.

We found that participant's behavior was close to rational, discounting distant goals at nearly the rate expected based on a rational calculation of cumulative risk. For example, participants valued a goal 5 steps away at close to the expected 59% discount rate. Subjective success probability ratings collected at the end of the experiment also reflected similar near-rational behavior. This behavior contrasts with classic findings of risk aversion in one-step gamble tasks with described or experienced risks. Using computational modeling, we found that exponential models were preferred to hyperbolic or probability distortion models. Previous research based on described risks has not yielded clear connections to real-world behavior. As decisions about learned risk are critical for balancing and pursuing goals, in an ongoing large-scale experiment using this paradigm we are testing whether variability in risky goal-directed decisions relates to altered decision-making associated with depression and anxiety symptoms.

P3-E-12 - Association between brain structure and genetic predisposition for high impulsivity in children.

Mari Shishikura¹, Filip Morys¹, Alain Dagher¹

¹McGill University

Introduction: Impulsivity is a complex heritable trait that is associated with many health problems, such as obesity and substance use disorders. Recent genome-wide association studies have identified genetic variants related to these phenotypes. However, it is unclear how these genetic variants affect brain development and lead to behavioral traits.

Here, we aimed to identify brain structures in children that are associated with genetic predisposition for adulthood impulsivity.

Methods: We included neuroimaging and genotyping data from 9 – 10-year-olds from the Adolescent Brain Cognitive Development study (N=4282). Following brain measures were included in our analysis: cortical volume, surface area, and thickness, subcortical volume, cellular density, and white matter integrity. We calculated 17 polygenic risk scores (PRSs), representing genetic risk of 17 different impulsivity traits, for each individual based on their genotype and published GWAS summary statistics.

In order to identify brain structural patterns that associate with genetic risk for impulsivity, we performed partial least square analysis between brain data and PRS data. Age, sex, and the first 10 genetic principal components were corrected in the analysis.

Results: The PLS analysis revealed two statistically significant latent factors relating PRSs and brain structures. The first and second latent factor accounted for 25.7% and 14.7% of the total covariance, respectively. The first latent factor showed that higher cortical surface area across the cortex was associated with higher PRSs for high-speed driving ($R=0.051$, 95% CI [0.021, 0.081]), drinking ($R=0.037$, 95% CI [0.007, 0.067]), and sensation-seeking ($R=0.031$, 95% CI [0.001, 0.060]), and lower PRSs for Barratt Impulsiveness scale ($R=-0.072$,

95% CI [-0.102, -0.042]) and delay discounting ($R=-0.031$, 95% CI [-0.060, -0.002]). The second latent factor showed that higher cortical thickness across the cortex associated with higher PRS for non-planning ($R=0.037$, 95% CI [0.006, 0.069]), and lower PRSs for sensation seeking ($R=-0.035$, 95% CI [-0.064, -0.005]), positive ($R=-0.046$, 95% CI [-0.076, -0.016]), and negative urgency ($R=-0.060$, 95% CI [-0.089, -0.030]).

Conclusion: We show that genetic predisposition for adulthood impulsivity is manifested in the early stages of brain development. We found that different impulsive traits have distinct genetic associations with cortical surface area and thickness across the cortex.

P3-E-13 - Not all discounts are created equal: regional activity and brain networks in temporal and effort discounting

Mohith Varma¹, Shanshan Zhen², Rongjun Yu¹

¹Hong Kong Baptist University, ²City University of Hong Kong

Costs like time delay and effort investment are generally discounted in decision-making. Standard economic models predict different types of costs are devalued in a similar manner. However, our review of rodent lesion studies indicated partial dissociations between brain regions supporting temporal- and effort-based decision-making. Another debate is whether options involving low and high costs are processed in different brain substrates (dual-system) or in the same regions (single-system). This research addressed these issues by using coordinate-based, connectivity-based, and activation network-based meta-analyses to identify overlapping and separable neural systems supporting temporal (39 studies) and effort (20 studies) discounting. Coordinate-based activation likelihood estimation and resting-state connectivity analyses showed immediate small reward and delayed large reward choices engaged distinct regions with distinct connectivity profiles but their activation network mapping engaged default mode network. For effort discounting, salience and sensorimotor networks supported low-effort choices, while the frontoparietal network supported high-effort choices. There was little overlap between temporal and effort networks. Our findings highlight the need to differentiate different types of costs in decision-making and understand discounting at both regional and network levels.

P3-E-14 - The interplay between stress, imagining the future, and intertemporal choice

Seoyoung Lee¹, Anita Restrepo, Hayoung Ahn, Greg Norman, Akram Bakkour

¹University of Chicago

Objective: Under acute stress, individuals seem to discount the value of future outcomes more than they would under no stress. However, imagining future events (i.e., prospection) has been found to reduce discounting of future rewards. Prospection is closely associated with memory because the possibilities we conjure may be limited by previous experiences. Studies have found that acute stress has negative effects on memory retrieval but remembering positive events may buffer the stress response. To clarify the interactions between stress, memory, prospection, and decision-making, this study investigates how stress might affect an individual's ability to prospect about the future which in turn might influence future-oriented decisions. We hypothesized (1) stress might affect prospection, and this might dampen the benefit of prospection on intertemporal choice (i.e., choosing between a smaller immediate reward or a larger future reward) or (2) stress might affect intertemporal choice, but the benefit of prospection might counteract the negative effect of stress.

Methods: Participants took part in a two-session study; each completing a stress and no-stress session. We induced acute stress using the Trier Social Stress Test while a control test was used for the no-stress session. Participants then completed a modified intertemporal choice task: they either imagined spending a certain amount of money in the future or described a scene prior to each intertemporal choice. The design resulted in four within-subject conditions (stress-imagine, stress-describe, no stress-imagine, no stress-describe). The standard hyperbolic discounting function was fit to choice data of each condition to obtain four discounting parameters per participant.

Results: Preliminary results ($N=12$) show that when under stress, prospecting about the future encourages participants to discount the future less compared to no prospection and to the same extent as under no stress. These results are based on behavioral data only, but we also collected participants' physiological stress responses throughout the study. The physiological data will be used to determine the efficacy of the stress manipulation and whether prospection reduces the stress response.

Conclusion: The present study investigates the interplay between stress, imagination, and decision-making to better understand how stress affects these interacting cognitive functions and how they might be resilient to the adverse effects of stress.

P3-F-15 - Information Seeking in Social Decision-Making is Driven by Instrumental Benefit

Alexandra Ramirez¹, Kenji Kobayashi¹, Joe Kable¹

¹University of Pennsylvania

Objective: To make effective decisions, people must seek and effectively weigh information that is relevant to the decision-making process. It has been established that, in nonsocial domains, people adaptively seek information based on its instrumental benefit for value-based decision-making. Little is known about what drives information seeking about others in social domains; however, research suggests that trait perceptions based on others' group membership, such as warmth and competence, drive social decision-making. Here we examined how people seek information about others in service of social decision-making and how it compares to nonsocial information seeking.

Methods: Participants ($N = 40$) played a modified trust game (social) and a lottery game (nonsocial). Across both games, they were asked whether to purchase information that reduced the associated uncertainty of a given choice. In the trust game, in some trials, participants were presented with an anonymous partner (undergraduate student)'s career aspiration and decided whether to trust them. In other trials, they were shown two possible partners, each with a different career aspiration, and asked whether they wanted to pay to reveal who was their true partner before choosing whether to trust them. The career aspirations had previously been rated on perceived warmth and competence. In the lottery game, in some trials, they were presented with a lottery and its odds of winning and were asked to choose whether to play. In other trials, participants were shown two possible lotteries, each with different odds of winning, and were asked to choose whether to pay to reveal the true lottery before choosing whether to play.

Results: Participants were more likely to trust partners with career aspirations rated higher in warmth (Kendall's $\tau = .71$, $p < .001$); there was no influence of perceived competence ($\tau = .22$, $p > .05$). Moreover, they were more likely to purchase information when the differences in warmth between the two potential partners was greater ($\tau = .46$, $p < .001$); again, there was no influence of the differences in competence ($\tau = .06$, $p > .05$). In the lottery game, participants were more likely to bet on lotteries with higher odds of winning ($\tau = .71$,

$p < .001$) and more likely to buy information when the winning-probability differences between lotteries was greater ($\tau = .65$, $p < .001$).

Conclusion: Our results demonstrate that people are capable of computing the value of information adaptively based on its instrumental benefit for upcoming decisions in both social and nonsocial contexts.

P3-F-16 - Social construction of value explains partisan alignment in behavior and neural representation

Benjamin Lu¹, Bo Shen¹, Jay Van Bavel¹, Paul Glimcher¹

¹New York University

Objective: It is widely accepted that partisan affiliations have a significant impact on the political preferences of individuals. Many US voters, for example, show strong political preferences over important issues that appear to be more tied to their political affiliations than to their direct self-interest. We recently developed a behavioral task designed to measure the strength and structure of how partisan affiliations influence preferences and employed fMRI scans to understand the neural processing.

Methods: Subjects viewed short synopses of real proposed laws that have appeared before US Congress. They were additionally asked to report their preference for the law based on the synopsis. In later blocks, they viewed and evaluated the laws again after learning the true percent of members of Congress in each political party who voted in favor of the law. This allowed us to search the neural data not only for the loci at which subjects' political preferences were represented, but also for neural activations that correlated with the degree of partisanship expressed by individual subjects.

Results: Our behavioral data showed that for many individuals, on both the left and the right, learning whether their political party supported a given law could rapidly and significantly shift their preferences for that law. Similarly, knowing whether the opposing political party supported that law could also strongly impact preferences. In addition to partisanship, the degree and direction of these shifts were notably influenced by a subject's psychological trait measures related to social and interpersonal relationships, such as desire for social esteem and approval of social hierarchies.

We examined the neural activations associated with political preferences and the rapid partisan-evoked shift in political preferences. We found areas that track political preferences in known economic value-related areas including ventromedial prefrontal cortex. We further found that these preference ratings employed the social cognition network of the brain involving temporoparietal junction and cingulate cortex. Finally, preliminary data suggest that dorsolateral prefrontal cortex is engaged when participants encounter unexpectedly high out-partisan support, which may point to the role of cognitive control in political judgments.

Conclusion: These findings indicate that political preferences use the common currency network for economic value and hint at possible policy solutions for the problem of polarization in American politics.

P3-F-17 - Reduced confidence in situations of observational learning

David Aguilar-Lleyda¹, Wojciech Zajkowski¹, Rei Akaishi¹

¹RIKEN Institute

Goal: agents learn about the value of the alternatives by choosing them and experiencing their reward, but often also by observing their peers making choices. Learning leads to being more confident in judging which alternative has the highest value, and in turn to choosing that alternative more. However, little is known on how learning from observation impacts choices between alternatives and confidence in their value, compared to learning from own choices. We investigated this question in three two-armed bandit experiments.

Method on each trial, one of two decks had to be chosen. Within each block, there was a high value deck that gave higher mean reward than a low value deck. After some trials, participants judged which was the high value deck, and confidence was rated on that judgment. In Experiment 1, reported in last year's meeting, we let participants decide whether, on each trial, they made the choice or observed another (simulated) agent choosing. Reward was only kept after own choices, but to create situations where observation was useful, we manipulated two variables across blocks. First, general reward could either be high or low, with the low value deck in low reward blocks being very likely to give negative reward. Second, participants could either observe agents who had more trials of experience than them, or the same amount of experience. These manipulations respectively captured observation as a way to avoid negative reward, and as a way to benefit from somebody else's experience.

Results: observation was more frequent in experienced than inexperienced observee blocks, and greater in low than high reward blocks. Low reward blocks were also accompanied by more choices of the high value deck, and by more correct judgments on the high value deck. However, absolute confidence was higher in high reward blocks. This dissociation could be due to a reward bias, where lower reward situations reduce confidence. Alternatively, they could be produced by an agency bias, where making more choices increases confidence. To disentangle between both hypotheses, we conducted Experiment 2, where the task dissociated lower reward from more observation. Confidence was still higher for high reward blocks, but also in no observation blocks. Results were replicated in Experiment 3, where reward feedback was given for both chosen and non-chosen decks.

Conclusion: compared to learning via own choices, learning by observing may reduce our confidence in the value of the alternatives. First, because more observation is coupled to situations of lower reward, and lower reward reduces our confidence. Second, because we feel more confident when making our own choices.

P3-F-18 - Verbal information exchange enhances collective performance through increasing group identification

Enhui Xie

Information exchange is a critical aspect of group decision-making that involves the conversion of individual private information into shared information known to all. Such information exchanges can take two primary forms which are verbal and nonverbal. Verbal information exchange involves discussion to exchange information, while nonverbal information exchange involves texting to exchange information. In the current study, we investigated whether and how information exchange enhanced collective performance by combining behavioral and neuroimaging approaches from the perspective of multiparticipant neuroscience. To evaluate collective performance, we measured the collaborative problem-solving abilities of triads working on a murder mystery case. The effect of different information exchanges (verbal and nonverbal) on collective success was first investigated in a behavioral study (Experiment 1, $N = 432$). Subsequently, the mediating role of group identification in this effect was examined in a functional near-infrared spectroscopy (fNIRS) study (Experiment 2, $N = 360$). We first found that verbal information exchange significantly enhanced collective performance

compared to nonverbal exchange ($F = 65.12$, $p < 0.001$). Moreover, both group sharing and group discussion positively contributed to this effect ($F = 7.92$, $p = 0.006$), with group discussion being more essential ($R^2 = 0.29$, $SE = 3.90$). Importantly, group identification mediated the positive effect of verbal information exchange on collective performance ($\beta_a = 0.56$, $SE = 0.27$, $t = 9.08$, $p < 0.001$; $\beta_b = 0.31$, $SE = 0.06$, $t = 5.58$, $p < 0.001$; $\beta_c = 0.40$, $SE = 0.03$, $t = 6.06$, $p < 0.001$). This mediation was supported by higher interactive frequency ($\beta_{a1} = 0.17$, $SE = 0.03$, $t = 2.81$, $p = 0.026$; $\beta_{a3} = 0.24$, $SE = 0.03$, $t = 4.08$, $p = 0.009$; $\beta_{b2} = 0.18$, $SE = 0.04$, $t = 3.36$, $p = 0.017$) and enhanced within-group neural synchronization (GNS) in the dorsolateral prefrontal cortex (DLPFC). These findings suggest that engaging in deeper exchange (such as information discussion) was a reliable predictor of group identification. We provided a multiparticipant theoretical model to explain how verbal information exchange enhanced collective performance. Our findings and theoretical model offer valuable insight into the decision-making process.

P3-F-19 - Network position biases mental representations of real-world social networks

Isabella Aslarus¹, Oriel Feldmanhall¹

¹Brown University

Humans are embedded in complex networks of social relationships. While the ability to represent and navigate networks has consequences for mental and emotional wellbeing, the manner in which people construct representations of their social networks is poorly understood. This is largely because the lion's share of past work has probed how people learn about small artificial social networks in the lab, which bear limited resemblance to real-world social networks that are large, complex, and constantly changing. In such studies, participants act as often-omniscient outside observers, but in the real world, your position in your network likely governs the information to which you have access, which in turn might influence the fidelity of your network representation. For instance, someone with many friends might receive information more often, and someone whose friends belong to distinct groups might receive more novel (less redundant) information. To explore how network position influences representation, we follow a social network of college freshmen ($N=185$) over the course of their first academic year. At multiple timepoints, we assess how a subset of network members ($N=100$) represent their network by probing their inferences about others' relationships. To our knowledge, this work is some of the first to longitudinally track a network from its inception, allowing us to link evolving network structure and changes in network members' positions with individual differences in representation. We find that network representations depend upon an individual's position within their network, operationalized with brokerage (the degree to which one bridges otherwise-unconnected others) and influence (the degree to which one is well-connected to well-connected others). Brokers and influencers both represent their networks more accurately. As a network member's brokerage increases or decreases over time, their representational accuracy increases or decreases correspondingly, an effect that does not extend to influencers. These effects are driven in part by brokers' and influencers' ability to identify latent network topologies, such as distinct communities, which they use to guide their inferences about others' friendships. In contrast, non-brokers and non-influencers exhibit representations dominated by observable, local features, such as individual relationships. Our results provide novel insights into the format of real-world social network representations, which form a basis for understanding naturalistic social navigation and behavior.

P3-H-20 - Meta-learning environment-specific learning rates

Jonas Simoens¹, Tom Verguts¹, Senne Braem¹

¹Ghent University

People often have to switch back and forth between different environments that come with different volatilities. While some environments require fast learning (i.e., high learning rates), others call for lower learning rates. Previous reinforcement learning studies have shown that people adapt their learning rates to their environment when differences in these statistics are clustered in time. However, these differences in learning rates could reflect emergent properties of participants' (non-environment-specific) responses to locally experienced reward prediction errors. As such, it remains unclear whether people can actually learn about environment-specific learning rates, associate them to relevant contextual features, and instantaneously retrieve them when revisiting environments (i.e., meta-learn the learning rate). Here, we hypothesised that people can learn to use and associate different, optimal learning rates to different environments. We used optimality simulations and Bayesian hierarchical modelling to demonstrate, across three experiments ($n = 273$), that people can alternate using two different learning rates, on a trial-by-trial basis, when switching back and forth between two two-armed bandit tasks in two different environments (i.e., casinos) that differ in volatility but are otherwise interchangeable. Results from a test phase suggest that participants learned to attribute these different learning rates to their respective environments, as there remained a small difference in learning rate between environments when the volatility was in fact similar. However, this difference was small because participants rapidly adapt to these newly experienced volatilities. Therefore, we ran another study that allowed us to estimate learning rates on a trial-by-trial basis, and thus to test whether differences in learning rates were already present in the first trial. On different locations, participants ($n = 50$) had to perform a continuous estimation task in 10-trial blocks on different locations around an island, with measurement uncertainties (and thus, optimal learning rates) depending on location. We again fitted several computational models to these data, and the model that fitted the data best had location-specific learning rates (especially in early trials), which differed across locations on the first trial (posterior probabilities of location differences all $> .97$), indicating that participants indeed associated learning rates to locations. We conclude that humans can learn to associate different learning rates to different environments, offering important insights for developing theories of meta-learning.

P3-H-21 - How the statistical information of others' donation influences individual donation behavior

Tao Jin¹, Iris M Vilares²

¹University of Minnesota, ²University of Minnesota, Twin Cities

People commonly engage in prosocial behaviors, actions intended to help or benefit others, which are susceptible to peer influence. Previous research on prosocial conformity primarily focuses on a single piece of social information (e.g., the averaged donation, the majority donation of a group). Instead, in an era of information explosion, we are often sequentially exposed to a large amount of diverse social information piece by piece. However, despite its prevalence, it remains unclear how a sequence of social information derived from different exemplars changes individuals' prosocial decisions, especially charitable donations. The current study investigated the effect of the mean (Low vs. High) and standard deviation (Low vs. High) of social information on individual donations. Participants ($n=317$) were randomly assigned to one of four conditions. They completed a charitable-donation task at first. After observing five others' donations, they had a chance to adjust their initial donations. We found that participants observing a high average

donation amount from others increased their averaged donations, whereas those observing small average donations decreased their own donations. However, we found no effect of the spread of others' donations on individual averaged donation shifts. In addition, the dispersion of individual donation amounts from the group norm decreased significantly after encountering social information in all groups. More interestingly, personal donations converged to the group norm much more after being confronted with others' donations with low variability compared to high variability. In terms of individual differences, we found that participant's initial donations were positively correlated with their empathy level (as measured with the Questionnaire of Cognitive and Affective Empathy) and negatively with psychopathy traits (from the Levenson's Self-Report Psychopathy Scale). These findings deepen our understanding of how both individual characteristics and the observed statistical information of others' donations can influence individual donation behavior, providing evidence for the effectiveness of the average level of social information in shifting individual prosociality, and the role of the diversity of opinions on individual convergence toward the group norm.

P3-H-22 - Forget and Forgive: a Memory-Based Mechanism for Increased Cooperation During Dynamic Group Formation

Wojciech Zajkowski¹, Ryan Badman², Masahiko Haruno³, Rei Akaishi¹

¹RIKEN Institute, ²Icahn School of Medicine, ³Center for Information and Neural Networks

Goal: Cooperation within groups is the cornerstone of human civilization. While much work has been done to explain how inherently selfish agents may sustain mutually beneficial cooperative relations, less is known about the influence of group formation dynamics on cooperation. Here, we test how such dynamics influence individual behavior, including choice strategies and overall levels of cooperativeness.

Method: We use a repeated, dyadic prisoner's dilemma (PD) task, where group size and composition can change on a trial-by-trial basis. The task consists of 180 trials, each being one instance of the PD game. Participants (N=87; 26 fMRI) begin the task by interacting with a single agent. At the start of a new trial, there's a 10% chance of being introduced to a new partner, who becomes a group member (max. group size = 5). At the end of 20% of trials participants can exclude the current partner.

Results: A logistic mixed model analysis revealed that group size and interaction distance (the number of trials since the previous interaction with the current partner took place) predicted increased cooperativeness ($\beta_{GS} = 0.12$, 95% CI [0.09, 0.15], $p < 0.001$; $\beta_{ID} = 0.05$, 95% CI [0.08, 0.02], $p < 0.001$) and decreased reciprocity ($\beta_{GS} = 0.04$, 95% CI [0.01, 0.06], $p < 0.001$; $\beta_{ID} = 0.05$, 95% CI [0.02, 0.08], $p < 0.001$). Hierarchical Bayesian model-based analysis revealed that this pattern of results can be best explained by a 'forgetting' model, where the value of reciprocity is moderated by the decaying memory of the partner's previous choice (>10 point difference in LOO and WAIC scores between best and second-best models). Whole-brain fMRI analysis ($p < 0.001$, cluster-wise FWE correction) revealed areas related to values of cooperation (Nucleus Accumbens (NAcc)), forgiveness (widespread activity across the fronto-parietal network and dorsomedial PFC), betrayal (dorsolateral prefrontal cortex (DLPFC)) and memory retention (precuneus and fusiform gyri (FFG)). Additionally, connectivity between the FFG and precuneus ($p = 0.01$), as well as precuneus and NAcc ($p < 0.01$), predicted memory retention.

Conclusion: Group size increase in a dynamic task setting leads to an increase in cooperation and a decrease in reciprocity. These effects are best explained by cooperation-reciprocity value tradeoff, where value of reciprocity is moderated by how well we can remember our partner's previous action. The integration of relevant memory-based information with subjective value is facilitated by the connectivity between FFG, precuneus, and NAcc.

P3-I-24 - Neural representations of attribute-based valuation across contexts.

Aniek Fransen¹, Kiyohito Iigaya², John O'Doherty¹

¹California Institute of Technology, ²Columbia University

Previous research has shown neural representations of stimulus attributes are integrated to compute an overall value judgment. As our goals and needs change we benefit from modulating the value of relevant items, necessitating flexible rather than static value representations. However, it remains unclear how our brain implements this flexible value modulation on an attribute level for complex objects. One hypothesis is that the value of attributes adapt to changing goals while the integration into an overall value judgment remains stable. An alternative possibility is that flexible value computation is accomplished by adapting the integration mechanism while attribute representations remain stable.

To differentiate between these hypotheses, we designed a paradigm that probes participant's (N=35) value judgments on a set of 75 unique clothing items across 3 goal contexts, while we obtained functional MRI data. To extract attributes of these high-dimensional stimuli we collected participant ratings for subjective attributes (e.g. comfort) and used methods from computer vision for objective attributes (e.g. contrast). Behavioral modeling shows that the value of attributes are adapted to the context and summed into an overall value judgment. As expected by the prediction that the integrated value of a stimulus is flexible to changing contexts by virtue of adapting individual attribute values, we find that the brain encodes both attribute-based and integrated value in ventral and dorsal medial frontal cortex, respectively. These results suggest that changes in the value of individual attributes underlie flexible value computation.

P3-I-25 - Memorability biases early memory sampling during risky choice

Hyuna Cho¹, Cendri Hutcherson¹

¹University of Toronto

Should we meet our friend for dinner this Friday? Should we buy groceries before or after work? Many real-world decisions require recalling information that is not immediately available in the environment. This sequential sampling of memory during choice has been theorized (Shadlen & Shohamy, 2016), but remains understudied. Here, we address an important question: whether and how we prioritize what information to sample from memory.

Participants (N=36) completed a 3 stage learning and memory task. First, in the value-learning task, 4 faces and 4 scenes were associated with 4 values (-5, -2, 2, 5). Second, in the cue-learning task, 16 unique face-scene pairs were associated with 16 object images, where some pairs were value-ambiguous (e.g. +2, -5) or value-congruent (e.g. +5, +5). Third, in the memory-based choice task, each object represented a gamble, where accepting meant a 50% chance of obtaining the associated face or scene value, which must be retrieved from memory. Critically, each block could multiply either face or scene values by x2 or x3, increasing the importance of the

cued attribute. We hypothesized that more important or more valuable information should be recalled earlier in the memory sampling process.

To demonstrate this, we show that a drift diffusion model with, rather than without, an early memory accumulation process better fits choices (DIC=191). Moreover, this model suggests that early face-, but not scene-, related sampling leads to more cue-biased choices ($\delta\mu_{\text{earlyFace}}=3.30$, $p<.001$) and more correct choices ($\delta\mu_{\text{earlyFace}}=3.10$, $p<.001$), suggesting an asymmetric relationship between attribute sampling and errors. However, this early sampling was not enhanced by the presence of external cues. Behaviourally, we found that faces were better remembered ($t(28)\text{paired}=0.286$, $p=.003$), and asymmetrically utilized during ambiguous trials (74.3%), suggesting a unique role of memorability in early memory sampling. Moreover, this asymmetry is better captured by simulations which model a prioritized memory-retrieval process, in which some types of information are sampled from memory first, rather than a static, evidence accumulation process, suggesting that retrieval-dependent decisions may be better captured by leveraging memory retrieval dynamics.

We demonstrate how memorability biases early choice dynamics, which may be uniquely important under uncertainty. We present preliminary work that suggests that memory retrieval dynamics can improve modeling of unique decision behaviors when choice attributes are recalled from memory.

P3-I-26 - Medial prefrontal-amygdala circuits contribute to cue-guided probabilistic reward decision-making in rats

Mudi Zhao¹, Shrishti Ramaiah¹, Stan Floresco¹

¹University of British Columbia

Decisions often carry uncertainty. An individual must assess the likelihood of a desired outcome and evaluate whether the potential gain outweighs the risk of failing to obtain it. Such outcome probabilities are commonly informed by extrinsic cues (e.g., using the cue of queues at a cafeteria to determine the odds of getting lunch in time for your next meeting). Previous research in rodents has implicated the basolateral amygdala (BLA) and the prelimbic and infralimbic subregions of the medial prefrontal cortex (PL/IL; homologous to areas 32/25 of the human cingulate cortex) in this type of decision-making. The PL and IL are functionally heterogeneous and exhibit distinct patterns of connectivity with the BLA. Here, we pharmacologically disconnected either PL↔BLA (n=14) or IL↔BLA (n=14) communication in rats while they performed on a cued probabilistic decision-making task known as the "Blackjack" task. Animals select between two levers: a safe/certain option that always delivers a smaller, 1-pellet reward, or a large/risky option that might deliver a 4-pellet reward. At the start of each trial, different auditory cues inform rats of whether the odds of obtaining the larger reward are good (50%) or poor (12.5%). The optimal strategy is to select the risky lever on good-odds trials, and the safe lever on poor-odds trials. Disconnection of prefrontal-amygdala pathways was achieved by a series of intracranial infusions of GABA receptor agonists, baclofen and muscimol. Disrupting either ipsilateral PL↔BLA or IL↔BLA signaling decreased risky choice on good-odds trials. The same effect was observed with disruption of contralateral IL↔BLA communication, but not contralateral PL↔BLA. When examining risky choice on poor-odds trials, disrupting either contralateral PL↔BLA or IL↔BLA signaling increased risky choice. Together, these findings confirm that communication between the BLA and both PL and IL mediates optimal decision-making on the Blackjack task. They also suggest dissociable contributions of these corticolimbic circuits and of ipsilateral vs. contralateral signaling. This might be attributed to differences in downstream and upstream connectivity patterns. BLA projections to the PL and IL are predominantly ipsilateral. In contrast, while IL sends ipsilateral projections to the BLA, PL sends both ipsilateral and contralateral projections downstream. Future research will directly target ascending vs. descending pathways and characterize direction-specific contributions of medial prefrontal-amygdala circuits in cued-guided probabilistic reward decision-making.

P3-I-27 - Specialized Neuronal Module in the Orbitofrontal Cortex for Economic Choice

Chengze Xu¹, Kuntan Ni¹, Xinying Cai¹

¹New York University Shanghai

A wide range of experimental and theoretical studies have established a clear connection between economic decision-making and subjective values represented in the orbitofrontal cortex (OFC). However, the specific function of the putative choice circuit in the OFC has rarely been examined under different choice tasks, leaving it unclear whether the OFC houses a functionally specialized module for economic choice or if it serves reward-based decisions in general.

To address this question, we conducted a study in which monkeys alternated between two distinct choice tasks in consecutive sessions: an economic choice task and a single attribute choice task. In the economic choice task, the animals had to choose between two juices of varying amounts (juice choice, JC), and their choice patterns reflected a subjective trade-off between juice taste and quantity. In contrast, the single attribute choice task required the animals to choose between the same juice in different quantities (quantity choice, QC), with the choice being directed deterministically towards the larger quantity offer, as reward quantity was the only differentiating attribute. These differing choice patterns indicated distinct mental operations under the two-choice tasks. The key question was whether the same population of neurons supported both JC and QC.

To investigate this, we recorded the activity of over two thousand neurons in two animals during both tasks. Our findings revealed that, during the JC task, OFC neurons encoded largely the same set of parameters as previously reported, namely offer value, chosen value, and chosen juice. Surprisingly, upon switching to the QC task, most of the neurons that were active during JC became inactive, while a largely complementary set of neurons encoded juice quantity in a spatial reference frame. However, the spatial choice signal was weak, suggesting that the OFC might not be primarily responsible for resolving QC.

To explore whether the representation of economic choice in the OFC, as observed in the JC task, is domain-general, we conducted a follow-up study where we trained one animal to switch between the JC task and an intertemporal choice (IC) task. In the IC task, the animal had to choose between an immediate juice reward in small amounts and a delayed juice reward in larger amounts. Neurons recorded from the OFC showed similar patterns of representation and were modulated by both tasks.

Taken together, our findings suggest the existence of a specialized neuronal module for economic choice in the OFC.

P3-I-28 - Transfer in searching for problem solutions: insights from a visual puzzle game

Yang-Fan Lu¹, Hang Zhang¹

¹Peking University

Objective: Problem-solving is a special case of sequential decisions (choosing a sequence of actions that leads to a desired goal), which provides a realistic task scenario for understanding flexible decision strategies such as analogical transfer. When the solver is faced with unknown action outcomes or combinatorial explosion of action sequences, one useful strategy is to apply the solution from an old problem to the new one. However, such analogical transfer can lead to negative outcomes as the old solution may be inefficient (e.g., Einstellung mental set effect) and even erroneous for the target problem. Here, we investigated how prior experience would influence the process of solving a new problem, focusing on the transfer of an erroneous solution.

Methods: In a computer visual puzzle game (adapted from “Baba Is You”), participants were given eight minutes to solve each new problem. Each participant went through five problems: two tutorial problems, one helper problem (providing essential information for solving subsequent problems), one target problem, and one source problem (whose solution is erroneous for the target problem). Participants were randomly assigned to an exposure group who encountered source before target (N=126) and a control group with target before source (N=128). We recorded participants’ action sequences in each problem and extracted the time points of each erroneous solution attempt in solving the target problem. We compared the exposure and control groups in the percentage, first time, and number of repetitions of attempting the erroneous solution.

Results: Participants who had prior exposure to the source problem did not have a higher percentage to attempt the erroneous solution (both groups ~40%), but had their first attempt earlier (hazard ratio=1.58, Z=2.21, p=0.028) and more repetitions (Wilcoxon rank-sum: Z=1.88, p=0.030). We further classified the erroneous solutions into two subtypes—“horizontal” (exact replication of the source actions) and “vertical” (a more effortful variant of the former)—and found an interaction effect of participants’ previous experience on the attempt of the vertical solution ($\beta=-1.24$, $t(246)=-4.14$, $p<0.0001$). For participants who failed in the helper problem and thereby lacked the information necessary for the correct solution (exposure: N=14; control: N=13), exposure to the source -increased the “vertical” attempt from 0% to 57.1%.

Conclusion: When they lack knowledge for the correct solution, people may choose to transfer solutions from analogous problems that may be erroneous.

P3-I-29 - Diabetes alters economically dissociable decision-making algorithms depending on the salience of reward scarcity in the environment

Chinonso Nwakama¹, Romain Durand-De Cuttoli², Zainab Oketokoun³, Samantha Brown³, Jillian Haller⁴, Adriana Méndez³, Mohammad Jodeiri Farshbaf³, Y. Zoe Cho⁵, Sanjana Ahmed⁶, Jessica Ables³, Brian Sweis²

¹University of Minnesota, Mount Sinai, ²Icahn School of Medicine, ³Mount Sinai, ⁴University of Scranton, Mount Sinai, ⁵Barnard College & Mount Sinai, ⁶Macaulay Honors College

Individuals with diabetes mellitus are at significantly higher risk of developing depression and other psychiatric disorders. Although diabetes is primarily characterized by chronic hyperglycemia, it remains unclear how impaired insulin function, which is known to have direct effects on neural activity, regulates motivated behavior. We characterized value-based decision-making of an insulin-deficient diabetic mouse model on a complex naturalistic neuroeconomic foraging paradigm. 40 8-week old CB57BL/6J male mice were injected with either vehicle (VEH) or streptozotocin (STZ), an antibiotic that ablates insulin producing beta cells in the pancreas, to induce hyperglycemia. Mice were then tested longitudinally across two months on the “Restaurant Row” task during which they foraged daily for their primary source of food while on a limited time budget. Mice learned to make serial decisions accepting or rejecting reward offers as a function of cost (delays cued by tone pitch) and subjective value (flavors cued by unique spatial contexts). Mice were trained on two different schedules during which the economic landscape (i) drastically or (ii) gradually progressed into an increasingly reward-scarce environment. Overall, STZ-treated mice earned less food but shifted meal consumption patterns in complex ways based on the revealed preferences of various flavors. Vicarious trial and error behavior, a proxy of deliberation, revealed decreased decision conflict for less-preferred flavors in STZ-treated mice. These findings were divorced from individual differences in economic choice policies, which were uniquely modulated in STZ-treated mice depending on their prior training schedules. Interestingly, we found that groups of mice valued the passage of time differently based on the type of choice being made. During change-of-mind decisions, mice became sensitive to the magnitude of time spent waiting, or “sunk costs,” in altering the probability of earning a reward but only after transitioning into a reward-scarce environment - except STZ-treated mice trained on a gradual schedule, who surprisingly never developed sensitivity to sunk costs. Deliberative and re-evaluative choice algorithms, which have been previously shown to be processed in physically separable circuits, may be differentially perturbed in a mouse model of insulin-deficient diabetes. These findings suggest complex relationships between glycemic regulation, the contrast of realized scarcity of the environment, and different types of opportunity costs interact to influence dissociable decision-making systems and fundamentally distinct behavioral computations underlying unique aspects of reward value.

P3-I-30 - Does affect bias investment decisions among professional investors?

Golijeh Golarai¹, Jennifer Eberhardt¹, Hazel Markus¹, Brian Knutson¹

¹Stanford University

Professional asset allocators manage more than \$69 trillion dollars globally on behalf of governments, universities, foundations, and companies. To perform their fiduciary duties in providing high rates of returns for their beneficiaries, asset allocators pursue various investment opportunities including for-profit financial intermediaries that are managed by professional fund managers, impacting the flow of capital across global markets. Yet, little is known about asset allocators’ decision-making processes. Are asset allocators’ decisions primarily driven by rational consideration, or shaped by implicit biases and affective processes that could detract from rational choice? Does professional expertise shield asset allocators’ rational choices from affective processes that typically impact decision making among lay individuals?

We designed an investment task where at each trial participants viewed information about an investment fund and rated the fund’s potential for return. Funds varied along 4 distinct levels of past performance. Each fund was paired with a unique identity from 1 of 4 face types (Black/White x Neutral/Smiling). Participants rated each fund’s potential for return on a 4 point scale. We posited (1) higher

ratings for funds paired with smiling faces compared to neutral faces, and (2) White viewers' higher ratings for funds paired with White faces compared to funds paired with Black faces.

To date, data from university students ($n = 10$) and professional investors ($n = 5$) show funds' past performance level impacted ratings across both cohorts as expected ($F=19$, $P<0.001$, $n=15$), and interacted with factors of race and expression ($P_s < 0.008$). However, there were no main effects of race or expression. We also found a cohort effect (student vs. professional participants) in fund ratings, and interactions with race and expression ($P_s<0.001$).

Within each cohort, facial expression impacted ratings ($P_s<0.03$) and interacted with race of faces ($P_s<0.046$). Surprisingly, the effect of race of faces was evident only among the cohort of White professional investors ($F = 63$, $P = 0.004$, $n = 5$), but not among students.

These findings suggest that affect impacts financial decisions across a range of professional expertise. Paradoxically however, racial bias was especially pronounced among professional asset allocators but not college students, perhaps reflecting divergence of college educated youth vs. older professionals, as the latter's attitudes maybe shaped by long-term immersion in a racially homogenous professional milieu.

P3-I-31 - A neuro-temporal decoding investigation of multi-attribute decision making

Matthew Bachman¹, Azadeh HajiHosseini¹, Sophie Faessen¹, Cendri Hutcherson¹

¹University of Toronto

Objective: Decision making frequently requires the calculation, weighing, and combination of the individual attributes of each option before making a choice. Prior neural research has characterized the spatial distribution of multi-attribute decision making by applying machine learning algorithms to fMRI data. Notably, these studies have reported significant activation in perceptual regions, in addition to those more commonly involved in the calculation of subjective value. However, fMRI's low temporal resolution has limited our ability to disassociate between the relative timing of perceptual and value signals. Our study fills this gap by applying machine learning algorithms to high-temporal resolution EEG data collected during a multi-attribute decision making task.

Methods: To accomplish this goal, we collected data from 27 participants who learned to associate different faces and colors with varying amounts of positive and negative values. Participants then briefly saw combined face-color stimuli, which they were asked to accept or reject based on the summed value of the face and color attributes. These points were accumulated and later converted into bonus money, thereby incentivizing participants to accurately accept positive total values and reject negative ones. Behavioral analyses indicated that participants were able to incorporate the face and color values in order to accurately accept net positive options or reject negative ones. Linear support vector machine classifiers were trained on the EEG data during the choice period in 10-ms time bins to decode when it could significantly predict activation for each condition.

Results Our neural results indicated broad temporal differences in when each attribute could first be significantly decoded, as well as how long these representations were maintained. More specifically, the perceptual attributes for faces and colors could be significantly decoded above chance starting around 110ms, while significant activation for their associated values was found at a slightly later latency. However, significant decoding for the value attributes was maintained for a much longer period of time, even after the stimuli were removed from the screen. The integrated face-color value signal was represented relatively weaker and slower, but also tended to be maintained for a longer period of time.

Conclusions: Our findings suggest key temporal distinctions between the formation of perceptual and value signals in multi-attribute decision making. Perceptual signals are fast but relatively brief, while value signals are slower and persistent.

P3-I-32 - Redefining Attribute Weighting in Food and Social Choices: Evidence for Non-linear Models

Dongqi (Sherry) Bao¹, Gaia Lombardi¹, Da Li¹, Carolina Feher Da Silva², Todd Hare¹

¹University of Zürich, ²University of Surrey

Background: Empirical data from risky choices, intertemporal choices, and multialternative choices highlight numerous violations of independence axioms, leading researchers to propose models focused on between-option interactions. However, other types of binary multiattribute choices remain underexplored, with linear difference models persisting as the default approach.

Methods & Results: We examine an overlooked, but striking behavioral pattern in food ($n=92$) and social choices ($n=102$) where individuals weigh an attribute differently based on the size of the difference in that attribute across options (for example, lower weight on health when health differences are larger, compared to when health differences are smaller). We demonstrated that the conventional linear model, which assumes a constant weight across attribute differences, is insufficient to capture this behavior. To address this, we proposed a modified drift-diffusion model incorporating a non-linear transformation of attribute value differences, supporting attribute-wise comparisons and emphasizing attribute similarity in multiattribute decision-making.

Furthermore, our study highlights the impact of negative and positive attributes on decision-making and underscores the importance of considering subjective transformations of individual attribute values based on loss aversion in specific attributes.

We showed that a drift-diffusion model incorporating a non-linear transformation of attribute value differences outperforms other models in both dietary (ELPD difference: -231.7) and social choices (ELPD difference: -6010.8), and when there is a distinction between negativity or positivity of attributes, including loss aversion in specific attributes (taste) is necessary. These non-linear transformations were relatively homogeneous across individuals and did not differ as a function of the healthfulness or prosociality of participant's choices.

Additionally, we found evidence for non-linear attribute differences across multiple stimulus presentation formats in our food choice dataset. This experiment included three presentation formats: only attribute ratings, only images of food, and both ratings and images. The non-linear transformation of attribute differences and attribute negativities persists across representations.

Conclusions: Our study reveals the shortcomings of linear difference models in multiattribute decision-making. Our proposed DDM, incorporating non-linear transformations of attribute differences and loss aversion in attributes, provides a more accurate description of choice behavior, demonstrating consistent success across various stimulus formats and individual behaviors.

P-I-33 - Competition between stimulus- and action-based learning strategies accounts for monkeys' behavior during dynamic learning

Jae Hyung Woo¹, Kathryn Rothenhoefer², Ken-Ichiro Tsutsui³, Fabian Grabenhorst⁴, Vincent Costa², Wolfram Schultz⁵, Alireza Soltani¹

¹Dartmouth College, ²Oregon Health & Science University, ³Tohoku University, ⁴University of Oxford, ⁵University of Cambridge

Objective: Intrinsic uncertainty in the reward environment requires the brain to run multiple models simultaneously to predict outcomes based on preceding cues, stimuli, or actions and ultimately, determine the reliable model that should control behavior at a time. Here, to examine whether and how competing models of the world interact to influence decisions, we re-analyzed three datasets of monkeys performing different dynamic learning tasks.

Methods: In the first two experiments, macaque monkeys were trained to perform a stimulus-based reversal learning (Costa et al, 2016) and dynamic foraging (Tsutsui et al, 2016). The third experiment was similar to the first one but with additional uncertainty about the correct model (stimulus- vs. action-based) of the environment (Rothenhoefer et al, 2017; controls only).

To identify the mixture of stimulus- and action-based strategies, we utilized the entropy of reward-dependent strategies (ERDS), equal to conditional entropy of stay/switch strategy given reward outcome. This yields two parallel quantities (ERDS_{stimulus}, ERDS_{action}) for how the two learning strategies were adopted by the animals.

Next, to reveal mechanisms underlying competition between competing models of the environment, we devised reinforcement learning (RL) models with two parallel learning systems for stimulus and action. In the static version of the model, the weight parameter ω_V , capturing the relative contribution of the two systems was fixed for the block. In the dynamic model, ω_V was adjusted every trial by comparing the reward prediction error (RPE) of each system.

Results: We found a competitive interaction between the two learning strategies, reflected as negative correlations between ERDS_{stimulus} and ERDS_{action}. This was strongest for the third task that included uncertainty about correct rule of the environment (Spearman's $r = -.78$, $p < .001$). Importantly, negative correlations were also observed for the other two tasks which did not require any action-outcome learning ($r = -.145$, $p < .05$; $r = -.11$, $p < .001$), raising the possibility that monkeys learned about action even when it was irrelevant. Five-fold cross-validation of models revealed that competitive learning models better accounted for monkeys' choice data across all tasks. Further, the improvement in fit for the dynamic model was the greatest for the third task which required arbitration between competing models of the environment.

Conclusion: Together, our results illustrate that primates' brain utilizes multiple learning systems to predict outcomes in uncertain environment, and arbitration between these systems could rely on concurrent RPE.

P3-J-35 - Identifying individual differences in (un)conscious emotion processing and decision-making via hierarchical drift diffusion models

Julia Schröder¹, Ute Habel², Han-Gue Jo³, Lennard Herzberg², Marie Kranz², Franziska Walter², Lisa Wagels²

¹RWTH Aachen University, ²Medical Faculty, Uniklinik RWTH Aachen, ³Kunsan National University

Objective: Previous research has shown that emotional stimuli can elicit behavioral responses at an unconscious level [1-5] but the optimal duration for unconscious stimulus processing is debated. The first study aimed to determine optimal duration and individual variations in unconscious stimulus processing. The second study compared underlying neural mechanisms in healthy controls (HC) and patients with major depressive disorder (MDD).

Methods: 40 HC rated emotional faces (happy, sad, neutral) at subliminal presentation times (8.3ms, 16.7ms and 25ms). A hierarchical drift diffusion model was compiled to evaluate the optimal presentation duration by estimating decision-making parameters. A second task was conducted with 60 MDD and 60 HC while assessing simultaneous EEG/fMRI. Images were presented as primers for 16.7ms (unconscious) or 150ms (conscious), followed by a target. Participants rated the target as depicting a happy, sad, or neutral face. A general linear mixed model was used to assess individual and task-related influences on decision-making (reaction time, RT), including group, stimulus emotion, primer presentation time, and primer/target (in)congruency effects.

Results: Stimulus awareness was reported in 65% of 25ms trials, 36% of 16.7ms trials, and 2.5% of 8.3 ms trials. Emotion-dependent responses were reflected in decreased drift rates during sad trials. Negative posterior drift rates of group means were found during 16.7ms trials. In the second task, MDD showed significantly higher RT compared to HC ($p = 0.019$, $t = 2.369$). Trials with unconscious primer presentation led to a higher RT compared to conscious primer presentation ($p < 0.001$, $t = -4.284$). Happy targets induced a faster RT than neutral and sad targets in both groups [MDD: happy-neutral: $p < 0.001$, $t = -14.104$; happy-sad: $p < 0.001$, $t = -24.065$; HC: happy-neutral: $p < 0.001$, $t = -19.448$; happy-sad: $p < 0.001$, $t = -31.856$]. No group differences regarding brain activity or the N170, an electrophysiological potential for face processing were found.

Conclusion: Subjective ratings and negative drift rates during 16.7ms trials indicate less effective stimulus processing compared to 25ms, suggesting that 16.7ms is the optimal duration for subconscious processing. Despite previous findings [5], we observed reduced RT during happy targets in both groups, possibly due to the high salience of happy facial expressions [6]. However, we speculate that regulation of emotional cues may require more cognitive effort in MDD compared to HC, leading to generally increased RT. There is currently no supporting evidence for differences in neural processing of unconscious stimuli.

P3-J-36 - Conceptual representations of emotions predict prosocial behavior

Marc-Lluis Vives¹, Joseph Heffner², Oriel Feldmanhall²

¹Leiden University, ²Brown University

People vary in how they use emotion concepts to distinguish between emotional experiences, which may impact the relationship between emotions and prosocial decision-making. We leverage methods which precisely characterize individual differences in emotion representation and test how this variation predicts prosocial choices. Eighty-one undergraduates at Brown University completed an emotion similarity task followed by a donation task. The similarity task used the spatial arrangement method (SpAM) where participants rearranged multiple emotion words such that proximities indicate similarity. To elicit personally consequential donations, participants received a \$100 endowment and were informed that they would be asked to donate a portion of their endowment to specific children from the Save the Children charity. 20 trials contained a photo of a child and 20 trials only used a silhouette. We use individual difference scaling (INDSCAL) to calculate individual differences in emotion representation. As expected, individuals conceptualize

emotions in a 2D space with positive arousal and negative arousal dimensions. Individual differences in the weight on negative arousal dimension, but not positive arousal, predicts decisions to donate to children in need, even when controlling for the emotional experience itself. Together, these results offer new understanding about the role of emotion knowledge in social choices.

P3-J-37 - Pupillary dilation predicts self-reported arousal during monetary incentive anticipation and outcomes

Yan Yan¹, Tara Srirangarajan¹, Brian Knutson¹

¹Stanford University

Pupil diameter is directly influenced by noradrenergic activity in the locus coeruleus, and is considered to reflect arousal (e.g., induced by perception of uncertainty or exertion of effort). Tasks involving anticipation and receipt of monetary incentives have been shown to physically increase self-reported aroused affect. While pupil diameter increases during anticipation of uncertain monetary gains, its connection to self-reported arousal has not been characterized.

To examine pupillary dilation and self-reported aroused affect at the same timescale, we recorded pupillometry with a virtual reality headset as 30 participants completed a Monetary Incentive Delay task. During each task trial, participants first saw a monetary cue indicating potential gain or loss, waited during an anticipatory delay, responded to a transiently presented target to either gain or avoid losing the previously-cued amount of money, and were notified of the outcome. Participants were probed about their current experience of affective valence and arousal on four-point scales, either during the anticipation or the outcome stage of each trial.

Monetary cues of increasing magnitudes increased both pupil diameter and self-reported arousal. Tonic and phasic pupillary dilation each correlated positively with current arousal, but their combination predicts affect better than each measure alone (anticipation: $R^2 = 0.18$; outcome: $R^2 = 0.11$). Assuming that the arousal and the valence axes are orthogonal (90°), a data-driven analysis found that anticipatory pupil dilation was only robustly associated with subjective arousal, whereas outcome-related pupil dilation was also slightly associated with valence (15° positive for hits, 35° negative for misses). Together, these analyses confirm that pupillary dilation mainly aligns with self-reported arousal, whereas its relevance to valence may be biased by the evaluative context.

To examine the temporal specificity of this association, we conducted cluster-based permutation on the regression weights of phasic pupillary dilation predicting self-reported affect over time. Results confirmed temporally specific coupling to the anticipation and outcome phases (permutation tests all $p < .001$). During the outcome phase, missing the target induced a longer coupling between pupillometry and self-report than did hitting the target, possibly consistent with a longer impact of negative than positive outcomes. Overall, this study provides an initial proof-of-concept 'read-out' of affect from pupillometry during processing of monetary incentives.

P3-J-38 - Negative affect and the computations underlying food choice in bulimia nervosa

Blair Shevlin¹, Loren Gianini², Joanna Steinglass², Karin Foerde², Caitlin Lloyd², Laura Berner¹

¹Icahn School of Medicine, ²Columbia University

Negative affect is a common precursor to binge eating among individuals with bulimia nervosa (BN). However, the mechanism by which negative affect precipitates binge eating remains poorly understood. In this study, we sought to characterize how negative affect impacted the decision-making process among individuals with BN making food choices.

Data were collected from 25 individuals who met DSM-5 diagnostic criteria for BN and 21 healthy controls (HC), group-matched for age and BMI. Participants completed a food choice task in two counterbalanced states: after a neutral affect induction, and after a negative affect induction. During the first phases of the task, participants used a 5-point scale to rate the healthiness and tastiness of 43 food items, each categorized as either low-fat or high-fat. In the subsequent Choice phase of the task, one food that was rated neutral for both health and taste was used as a Reference Item, and participants made choices between this Reference Item and the other food items. At the task's conclusion, participants were given a chosen food from one randomly selected trial.

We used a time-varying diffusion decision model (stDDM) to study decision dynamics during the Choice phase. The stDDM introduces a starting time parameter that indicates the delay with which different attributes enter the evidence accumulation process. The model was estimated using a hierarchical, Bayesian approach in JAGS.

We found that across affective states, the BN group put more weight on health than the HC group ($\mu = -1.25$, 95% HDI $[-0.40, -2.15]$), while the HC group put relatively more weight on taste ($\mu = 1.80$, $[0.46, 3.97]$). The HC group accumulated taste information before health information for both LF foods ($\mu = -0.42$, $[-0.64, -0.20]$) and HF foods ($\mu = -0.66$, $[-1.04, -0.30]$); negative affect had no effect. In contrast, the BN group accumulated taste information first only for HF foods ($\mu = -0.65$, $[-1.17, -0.23]$), and this start time bias toward taste information was exaggerated after the negative affect induction ($\mu = -0.29$, $[-0.56, -0.02]$).

Our results suggest that negative affect may influence food choices in BN by changing the latency with which certain information enters the decision process. Although health information is strongly weighted by individuals with BN, negative affect delays its entry into the evidence accumulation process. With this longer delay, taste information has more time to bias decision makers towards high-fat food choices before health information can come online. These dynamics may explain why binge eating is more likely during periods of high negative affect, but dietary restriction is more likely during periods of low negative affect.

P3-J-39 - Shared neural representations and temporal segmentation of political content predict ideological similarity

Daantje De Bruin¹, Jeroen van Baar¹, Oriël Feldmanhall¹

¹Brown University

Despite receiving the same sensory input, opposing partisans often interpret political content in disparate ways. By jointly analyzing controlled and naturalistic fMRI data ($N=44$) and by leveraging a combination of representational similarity analysis, inter-brain synchrony, and temporal event segmentation, we uncover the neurobiological mechanisms explaining how these divergent political viewpoints arise. Individuals who share a political ideology have more similar behavioral and neural representations of political words, experience greater neural synchrony during naturalistic political content, and temporally segment real-world information into the same meaningful units. In the striatum and amygdala, increasing inter-subject similarity in neural representations of political concepts during a word reading task predicts enhanced synchronization of BOLD time courses when viewing real-time, inflammatory political videos, revealing that polarization can arise from differences in the brain's affective valuations of political concepts. Ideological similarity was also predicted by more similar temporal event segmentation in the dmPFC, which appeared to reflect partisan differences in emotionally and semantically distinct experiences. Together, this research shows that political ideology is shaped by semantic representations of

political concepts processed in an environment free of any polarizing agenda, and that these representations bias how real-world political information is construed into a polarized perspective.

P3-J-40 - Social Decision Making: How stereotypes and group membership interact to affect fairness, a fMRI study

Harshil Vyas¹, Paloma Diaz-Gutierrez², Carolyn Declerck¹, Christophe Boone¹

¹University of Antwerp, ²Universiteit Antwerpen

Navigating social situations relies heavily on multiple mechanisms of social perception. Much research in social psychology has corroborated that people are more fair towards others who are close to them, as well as to people who are similar to them. Overt characteristics such as race and gender tend to shape expectations in interpersonal contexts, and so does categorization of self and others in groups. However, previous research studies that have examined the influence of closeness (in-group favoritism) and similarity (stereotypes) have done so individually, as well as interchangeably. We proposed instead that (RQ 1) categorizing others based on group membership is fundamentally different from categorizing others based on stereotypes, and that they are the consequence of different neural processes. We further proposed that (RQ 2) stereotyping can be attenuated or strengthened depending on group membership. Using a minimal group paradigm, participants ($n=40$, female= 28) in this study were asked to choose the team that they would like to belong to by choosing a painting from a set. Thereafter, they interacted with in- and out-group members from 20 different social categories in a hypothetical, modified Dictator Game. For each trial (120 decisions), participants were first shown the social category of their partner (characteristic/profession/social status/ethnicity - surgeon, British, & welfare recipient for example) with cues indicating their group membership. Subsequently, faced a decision screen wherein they were asked to choose between one fair and an unfair option (3 variants). Each trial was an advantageous inequity decision. Behavioral results and manipulation checks thus far suggest successful replication of findings from previous research in terms effect of warmth (understood through the two-dimensional framework of stereotype content model) on fairness, and of a strong in-group bias, highlighting the effectiveness of minimal groups. We observed that participant's giving behavior was significantly influenced by warmth, such that increase in perceived warmth led to an increase in fairness. Preliminary examination into our neural results too show pattern activation in areas associated with social information processing when participants were introduced to their interaction partners. Uncorrected neural data with a 0.001 threshold have shown thus far, activation in middle temporal gyrus, middle frontal gyrus, lateral PFC, angular gyrus, and cerebellum when making decisions for in-group members. Further whole and ROI analyses will soon be conducted on the acquired data to investigate the interaction of group membership and warmth on fairness.

P3-J-41 - Contributions of insula and superior temporal sulcus to interpersonal guilt in social decisions

Maria Gädeke¹, Tom Willems², Omar Salah Ahmed¹, Bernd Weber¹, René Hurlemann³, Johannes Schultz¹

¹University of Bonn, ²University of Bern, ³University of Oldenburg

This study investigated the neural mechanisms involved in feelings of guilt evoked by social decisions. In two studies (one inside fMRI), participants repeatedly chose between a safe and a risky monetary outcome. Across conditions, participants chose just for themselves (control condition), chose for both themselves and a partner (active social condition), or were recipients of the choices of their partner (passive social condition). If the risky option was chosen in a social condition, both participants could receive either the high or the low outcome of a lottery with 50% probability, independently of each other. Participants were shown the outcomes for themselves and their partner on each trial, and reported their momentary happiness every few trials. Participant happiness decreased following both low lottery outcomes for themselves and for the partner. Crucially, happiness decreases that followed low outcomes for the partner were larger when the participant rather than their partner had made the choice, which we interpret as guilt. Insula responses to low lottery outcomes for the partner were larger following participant's rather than partners' choices. Variations in happiness were well explained by temporal difference models relying on participant's and partners' rewards and reward prediction errors. An fMRI analysis based on these models revealed a cluster of voxels in left superior temporal sulcus whose activation fluctuated with reward prediction errors experienced by the partner, and this response varied depending on who made the choices leading to these prediction errors. Our findings suggest complementary roles of the anterior insula and the superior temporal sulcus in the neural mechanisms underlying interpersonal guilt.

P3-J-42 - An Integrative Model of Information Sharing Decisions: Semantic features, neural correlates, and affective experience

Saskia Balny D'avricourt¹, Bruce Doré¹, Laurette Dube¹

¹McGill University

Decisions to share information can lead to widespread diffusion of content and shape important societal outcomes. However, the underlying mechanisms that drive these decisions are not fully understood. Prior work has asked how different semantic features (specific words), and neural responses can influence sharing behavior. Research has shown that brain regions associated with self-related and social cognition indirectly influence population-level sharing by activating the brain's value system, including the ventral striatum and ventromedial prefrontal cortex. Moreover, studies found that emotional or expressive messages are more likely to be shared. However, much is still unknown about how these textual and neural features relate to and interact with core psychological processes driving sharing behavior. To address this, we develop an integrative framework, based on appraisal theories of emotion, to explore how content features, brain responses, and affective experiences together shape sharing behavior. Focusing on health-related messaging, we measured content features with text analysis, brain activity with neuroimaging (fMRI) ($n = 41$), and affective experience with subjective reports ($n = 247$). By applying dimension reduction techniques, we were able to identify a parsimonious set of underlying components and examine their distinct and overlapping contributions. Results indicated that present-tense and emotion-related content features (adj $R^2 = 0.11$, $p = .018$), experiences of information relevance and affective impact (adj $R^2 = 0.17$, $p = .001$) and reward-related brain responses (adj $R^2 = 0.16$, $p < .001$) were all predictive of article sharing. When combined, these sets of predictors each contributed incremental predictive validity (overall adj $R^2 = 0.44$, $p < .001$). Further analysis of the relationships between these components reveal a pattern of brain responses related to appraisals of self and social relevance (adj $R^2 = 0.36$, $p < .001$) which are involved in sharing decisions. When assessing brain responses to content features, reward-related responses demonstrated some relation to features such as personal pronoun use and social and emotional content (adj $R^2 = 0.05$, $p = 0.03$), which are prominent in highly shared articles. In summary, our study demonstrates that brain responses, subjective reports, and text features each provide unique predictive value for information sharing. This work sheds light on the psychological and neural mechanisms that drive information sharing, as well as the objective content features that tend to engage these mechanisms.

P3-L-43 - How Working Memory Overload May Reduce Dietary Self-control: A Neuroeconomic Experiment

Łukasz Tanajewski¹, Jakub Skałbania^{1,2}, Marcin Furtak³, Marek Wyppych³, Todd Hare⁴

¹Kozminski University, ²University of Warsaw, ³Nencki Institute of Experimental Biology, ⁴University of Zürich

Aims: Previous experiments have shown poorer dietary self-control under increased working memory load (IWML). Neuroeconomic studies on food choices revealed increased activity in dorsolateral prefrontal cortex (dlPFC) in dietary self-control, the integration of health and taste values in ventromedial prefrontal cortex (vmPFC), and that dlPFC supports the incorporation of health into vmPFC value signals. We studied the neural base of dietary self-control under IWML, and postulated that IWML reduces (H1) choice-related dlPFC activity, (H2) the functional connectivity between dlPFC and vmPFC, and (H3) the link between health value and vmPFC activity, with (H4) effects in H1-H3 resulting in impaired dietary self-control.

Methods: Participants (n=49) rated the healthiness and tastiness of foods. In the fMRI scanner, they made 90 food choices twice: after memorizing one and seven digits. Dietary self-control was defined as choosing a healthier over tastier food item. Choice-related brain activity controlled for load-specific activity was analyzed with generalized linear and psychophysiological interactions models. Family-wise error corrected for cluster size (FWEc) and region of interest approach (with two-sided nonparametric test p-values) were employed to test H1 and H2-H4, respectively.

Results: (1) IWML reduced bilateral dlPFC activity (FWEc; $p < .05$). (2) The functional connectivity between right dlPFC (seed: 6mm around peak activity) and vmPFC was reduced under IWML ($p = .04$). (3) Parametric modulation showed that the healthiness benefit from a healthy choice correlation with vmPFC activity was weaker under IWML ($p = .01$). (4) IWML did not change the probability of dietary self-control on average nor did we find group-level effects on vmPFC or dlPFC activity patterns. Using a measure of self-control which incorporated decision strength (Tresponse-1) and the imprecision of healthiness and tastiness ratings, we showed reduced dietary self-control under IWML (mixed model, $p = .02$), and the correlation ($r = .32$) between the negative effects of IWML on self-control and on left dlPFC activity ($p = .03$).

Conclusions: Shifting dlPFC resources from choice-related to load-specific activities and the resultant weaker dlPFC support of health values incorporation in the vmPFC signals might explain poorer dietary self-control under load found in previous experiments. Using a binary measure, we did not find an effect of memory load on self-control in food choice, nor brain-behavior interactions. More sensitive lab measures of dietary self-control are needed to study these effects, possibly due to noisier decisions under load.

Funding: The study was supported by Polish National Science Center Grant 'Cognitive resources and self-control of impulses: a neuroeconomic approach' (No. 2017/26/D/HS6/01157). PI: Łukasz Tanajewski.

P3-L-44 - Effect of transcranial direct current stimulation (tDCS) on FRN and P2 during the performance of auditory monetary incentive delay task

Anastasia Grigoreva¹, Aleksei Gorin¹, Valeriy Klyuchnikov¹, Anna Shestakova¹

¹HSE University

Transcranial electrical stimulation (tES) serves as a powerful technique for assessing the causal role of specific brain regions in behavior, including decision-making. While tES studies investigating sensorimotor function yield unequivocal results, cognitive research reveals considerable ambiguity and diversity in stimulation-related effects. The consequences of transcranial direct current stimulation (tDCS) on cognitive functioning are not consistently predictable based on the current direction (cathodal or anodal), which limits its applicability in cognitive research.

In the present study, we further explored tES effect ambiguity in cognitive tasks by modulating medial-frontal cortex (MFC) activity in an auditory monetary incentive delay (MID) task, where participants ($n = 30$) responded to acoustic cues encoding different expected monetary losses, using cathodal tDCS. We analyzed feedback-related negativity (FRN), reflecting a reward prediction error processing when participants encountered losses compared to no losses during two subsequent MID task sessions, and exogenous P2 response to stimulus onset unrelated to anticipated monetary contingencies. We anticipated an inhibitory cathodal tDCS effect on both P2 response and FRN.

Contrary to our hypothesis, we observed a facilitatory effect of cathodal tDCS on FRN ($F(2, 28) = 6.99$, $p = 0.013$, $\eta^2 p = 0.2$) along with suppressed learning effect of P2. The difference in P2 amplitude between the first and second sessions, observed in the sham group ($F(1, 13) = 11.182$, $p = 0.002$, $\eta^2 p = 0.493645$), was absent in the group that received cathodal tDCS. We discuss the controversial effects of tDCS on P2 and FRN during auditory MID performance in light of brain plasticity and reward-learning mechanisms, correspondingly. Importantly, our findings regarding the multidirectionality of tDCS effects on cognitive function challenge the utility of tES as a readily employable method for testing brain causality in highly complex neurocognitive events, such as decision-making.

P3-L-45 - Designing a sustainable economy with deep reinforcement learning

Raphael Koster¹

¹Google DeepMind

Successful resource allocation mechanisms maximize prosperity by encouraging recipients to make reciprocal contributions to the economy. How can we allocate resources in ways that foster sustainable exchange? We consider an iterated multi-player trust game, in which four players interact with a resource pool over 40 time steps, that is managed by an allocation mechanism. We use deep reinforcement learning (RL) to design an allocation mechanism that promotes sustainable contributions from human participants ($n=1316$) to a common resource pool, leading to a large surplus and an inclusive economy. The agent is trained by populating a simulation with neural networks representing players. These neural networks have been trained from human behavior on the same task (via 'behavioral cloning'), and show a strikingly similar behavior even across conditions not contained in the training data. The RL agent increases human surplus over baseline mechanisms based on unrestricted welfare or conditional cooperation, conditioning its generosity on available resources and temporarily sanctioning defecting players. Examining the agent policy allowed us to develop an explainable mechanism that performed similarly and was more popular among players. This mechanism was more restrictive (conditional cooperation) when the resource was depleted, but gave unrestricted access when the resource was close to the maximum. In conclusion, we show that behavioral cloning can be used to build faithful simulation of human economic behavior and deep reinforcement learning can be used to discover mechanisms that promote sustainable exchange.



Symposium 1

S1.1 - A common cause for multiple suboptimality in evidence accumulation

Diksha Gupta¹, Brian Depasquale², Charles Kopec³, Carlos Brody³

¹University College London, ²Boston University, ³Princeton University

Objectives: It has long been known that even on simple tasks, subjects with significant experience often deviate from the optimal policy. A prominent example of this shows up in perceptual decision-making tasks in which subjects are asked to report the category of a noisy stimulus, by accumulating evidence over time. In such tasks, subjects often assume dependencies between successive trials even if they are independent, and consequently show history biases in their decisions. In addition, on a sizeable fraction of trials subjects seem to choose stochastically, independent of the evidence - such errors are called lapses. While these two suboptimality (history biases and lapses) are routinely assumed to arise from distinct processes, there are several hints in the literature that suggest that they covary in their prevalence, and that their proposed neural substrates overlap. Here, we propose a common cognitive process that can parsimoniously explain the links between these suboptimality.

Methods and Results: We explore the idea that history biases reflect a misbelief about non-stationarity in the world, and demonstrate that normative evidence accumulation under such beliefs gives rise to choices that are both history-dependent and appear to be evidence-independent (i.e. akin to lapses). This process corresponds to an accumulation-to-bound model with history dependent updates to the initial state of the accumulator. We tested this model in a large rat dataset consisting of choices from 152 rats trained on an auditory decision-making task and confirmed its predictions using detailed model-fitting. Despite heterogeneity in history biases and lapse rates in this population, we show that a substantial fraction of lapses can be explained by the presence of history dependence in initial accumulator states. Finally, we tested the reaction time predictions of the model in a novel task in rats, and confirmed that the data showed signatures of initial state updating. The model could successfully capture both choices and reaction times, while simultaneously accounting for history biases and lapses.

Conclusions: Altogether, our results challenge the prevalent assumption that history biases and lapses have independent causes and instead show that they can arise from a common cognitive process. Our model improves the ability to precisely predict decision-making dynamics both within and across trials and renders choices that were previously thought to be stochastic (i.e. lapses), predictable.

S1.2 - Integrating information and reward into subjective value: humans, monkeys, and the lateral habenula

Ethan Bromberg-Martin¹, Yang-Yang Feng², Takaya Ogasawara², J. Kael White², Kaining Zhang², Ilya Monosov²

¹Washington University School of Medicine, ²Washington University

Objective: Humans and several animal species can be strongly motivated to seek information about uncertain future rewards. Remarkably, they seek information even when it has no objective value for controlling the outcome, suggesting that they endow information with subjective value of its own. There has been a recent explosion of research on this phenomenon, which has been called observing behavior in psychology and temporal resolution of uncertainty in economics. However, an unanswered question is whether the computations that assign subjective value to information are conserved between humans and other species. If so, we could leverage animal models to uncover the neuronal populations that are responsible for conserved information value computations and their causal influence on decisions.

Methods: We designed analogous information choice tasks for humans (n=565) and monkeys (n=4). Individuals choose between options with multiple attributes, including cues that are either informative or non-informative about future outcomes, and different probability distributions of rewards (money for humans, juice for monkeys). This let us measure and model the subjective value individuals assign to information; how they compute information value using reward uncertainty, expected reward, and other attributes; and integrate information and reward into the total value of an option. We investigated neuronal networks for these value computations by recording neurons in monkeys from connected areas with information-related activity, anterior/ventral pallidum (Pal, n=294) and lateral habenula (LHb, n=375).

Results: Value computations were remarkably similar in humans and monkeys. For example, both valued information based on the uncertainty it would resolve ($\beta > 0$, $p < 0.001$ in each animal; mean $\beta = 1.40 \pm 0.06$, $p < 0.001$ in humans), and were best fit as computing this "uncertainty" with a specific family of mathematical uncertainty measures ($\Delta BIC > 720$ vs competing models). Both LHb and Pal responded to all attributes needed for decisions, but only LHb neurons predominantly integrated many information- and reward-related attributes to reflect the total subjective value of options. Further, trial-to-trial fluctuations in LHb value signals predicted ongoing decisions (mean $R = +0.06$, $p < 0.001$), while electrical stimulation coincident with LHb value signals modified ongoing decisions (mean $\beta = 0.71 \pm 0.19$, $p = 0.0007$, $n=22$).

Conclusion: Our data uncover fundamental computational principles for evaluating information and reward and integrating them into total subjective value, and implicate LHb in these conserved value computations that guide online

S1.3 - How Multifaceted Motives impact Evidence Accumulation

Laura Globig¹, Tali Sharot¹

¹University College London

Objective: Several, sometimes conflicting, motives impact evidence accumulation. Consider an employee tasked with investigating their colleague for fraud. They would like to find their colleague 'not guilty'. But to do well in their job they have to give an accurate verdict. To reach a judgement they gather evidence until a certain degree of confidence is achieved. What this threshold is and when it is reached may depend on how noisy the evidence is, how much the employee cares about their job and whether the evidence is pointing towards the desired result. While such factors have been studied in isolation, in real life they often coexist. Here, we examined how these factors together influence evidence accumulation.

Methods: Participants (N=70) completed a random dot kinetogram task, in which they judged whether most dots were moving right or left. Across trials we varied: motion coherence (i.e., noise); whether participants were incentivized for accuracy (i.e., paid for correct judgements); and participants' preferred judgement. The latter was manipulated by telling participants that if the dots moved in a previously specified direction (right/left, counterbalanced) they would receive a bonus regardless of their response.

Results: Drift Diffusion Models in which coherence, preference, and accuracy incentives could modulate the drift rate (i.e., rate of accumulation), starting point, boundary separation and non-decision time were fit to the data. The winning model showed that accuracy incentives modulated boundary separation such that participants were more cautious when incentivized for accuracy (± 0.061 ; 95%CI

[0.022; 0.095]), thus taking longer to reach a judgment ($\text{std.}\hat{\tau}^2=0.014\pm0.005$, $p=0.005$). Drift rate, however, was modulated by coherence and preference, such that the rate of evidence accumulation was greater when coherence was high ($v=0.382$; 95%CI [0.296; 0.466]) and when the dots were moving in the preferred direction ($v=0.124$, trend: 95%CI [-0.005; 0.251]). Thus, judgments were faster ($\text{std.}\hat{\tau}^2=-0.053\pm0.009$, $p<0.001$) and more accurate ($\text{std.}\hat{\tau}^2=0.4\pm0.042$, $p<0.001$) when coherence was high; and by trend less accurate when the dots were moving in the undesired direction ($\text{std.}\hat{\tau}^2=0.137\pm0.08$, $p=0.079$). Healthy individuals were more biased towards preferred judgments than depressed individuals ($t(67)=2.058$, $p=0.044$, Cohen's $d=0.515$). No interactions were significant.

Conclusion: Distinct motives are integrated into the decision-making process by independently impacting different aspects of the evidence accumulation process. Thus, when people are motivationally biased, accuracy incentives and noise reduction can lead to improved judgments.

S1.4 - The informational cost and utility of expressing emotions: a computational model of social signalling

Yi Yang Teoh¹, Cendri Hutcherson¹

¹University of Toronto

Emotion expressions are vital signals for social interactions. How do people decide when to express, suppress, or exaggerate, what they feel in these social interactions? In an experimental study, we demonstrate that participants anticipate both social costs and social utility to (angry, joyful, and sad) expressions and strategically express their emotions as information signals to regulate their partner's social preferences. Research in economic games suggests that people often feel and express negative emotions like anger in the face of inequality, perhaps even to an exaggerated extent as a negotiation tactic. Simultaneously, a largely separate literature suggests that individuals sometimes suppress emotion expressions due to their potential social costs. We compared how participants felt and what they expressed when an anonymous partner shared different amounts of money with them ($n = 246$) and modelled their emotional responses as an evidence accumulation process within a hierarchical Bayesian framework using a novel multi-attribute Racing Diffusion Model to quantify these perceived costs and benefits to emotion expression. Corroborating extensive theorizing about the potential costs to expressions, we found that participants appeared to anticipate large costs to expressing anger which subsequently led them to suppress it (Cohen's $d = -0.498$, 90% HDI [-0.784, -0.244]). However, this was not true for all emotion expressions. Instead, we found that participants also anticipated significant social utility to expressing both joy (Cohen's $d = 0.310$ [0.019, 0.557]) and sadness (Cohen's $d = 0.342$ [0.069, 0.611]) during these economic interactions that led them to exaggerate these expressions even when participants did not self-report feeling them. Furthermore, our study suggests that these costs and benefits to expression were highly dynamic, with participants anticipating greater utility to suppressing anger that arose out of the losses they incurred (Cohen's $d = 0.343$ [0.083, 0.627]), and exaggerating/faking joyful expressions when their partners' chose less unequal distributions (Cohen's $d = -0.353$ [-0.610, -0.076]).

Our results here strongly suggest that emotion expressions constitute a value-based decision where people strategically weigh the anticipated social costs and benefits of each expressive option when deciding when to express, suppress, or exaggerate their emotions to social partners. Critically, these findings have strong implications for how we understand social behavior, underscoring the potential importance of affective processes in interpersonal communication.



Symposium 2

S2.1 - Beyond the Matrix: Experimental Approaches to Studying Social-Ecological Systems

Joel Leibo¹, Raphael Koster¹, Marco Janssen², Uri Hertz³

¹Google DeepMind, ²Arizona State University, ³University of Haifa

Note: this talk is intended as the opener for a symposium. So it's meant to be a broad overview, introducing the audience to the topic. The other three talks submitted as part of this symposium are all describing specific primary research studies.

Studying social-ecological systems, in which agents interact with each other and their environment is a challenging but important task. In such systems, the environment shapes the agents' experience and actions, and in turn collective action of agents changes social and physical aspects of the environment. Experimental and computational approaches to studying complex social behaviors and processes have come a long way since the 1950s. However, emphasis on directly mapping the paradigms that are most computationally convenient (matrix games) to their direct analogs in the laboratory may have impoverished experimental design. Modern artificial intelligence (AI) techniques provide new avenues to model complex social worlds, preserving more of their characteristics. These techniques can be fed back to the laboratory where they help to design experiments in more complex social situations without compromising their tractability for computational modeling. This novel approach can help researchers bring together insights from human cognition, sustainability, and AI, to tackle real world problems of social ecological systems such as climate change, pandemics, and conflict resolution.

S2.2 - Conditions for cooperation in Social-Ecological Systems

Marco Janssen¹

¹Arizona State University

The tragedy of the commons narrative proposed in 1968 by biologist Garrett Hardin adopts the simple logic that people who share a common resource will overharvest it. This led to the advocating of privatization or nationalization of shared resources. The Hardin paper is one of the most cited and influential academic papers, but it is disproven by Elinor Ostrom and her colleagues who did systematic empirical research of many cases all over the world. Communities can organize themselves in managing their shared resources under the right conditions.

Ostrom's observation emphasized the complexity of social-ecological systems that impact the outcomes. She proposed a behavioral theory of collective action that takes into account the complexity and diversity of human decision-making. Despite those observations, there has been a reluctance to embrace complexity due to the lack of formal approaches to study this complexity. The combination of MARL and novel behavioral experiments makes a more formal study of self-governance possible.

In this contribution, I will review different types of behavioral experiments in the lab and the field where ecological and social complexity is included to understand the conditions for cooperation in social-ecological systems. Important factors that explain cooperation are procedural justice, trust, and fairness. Moreover, new results will be presented where communication patterns in more than 1500 rounds in more than 100 groups in four resource types (groundwater extraction, surface water irrigation, spatial foraging, and community infrastructure) are analyzed.

S2.3 - Studying large-scale collaborations in online communities

Natalia Vélez¹

¹Princeton University

Humans have developed technological repertoires that have enabled us to survive in virtually every habitat on Earth. However, it can be difficult to trace how these technologies came to be—folk histories of technological achievement often highlight a few brilliant individuals, while losing sight of the rest of the community's contributions. In this talk, I will present work analyzing player behavior in One Hour One Life, a multiplayer online game where players can build technologically complex communities over many generations (N = 22,011 players, 2,700 communities, 428,255 lives lived, 127,768,267 social interactions detected). This dataset provides a unique opportunity to test how community-wide expertise and division of labor shape technological development: Players can form communities that endure for many generations, and they can combine thousands of unique materials to build vast technological repertoires. Overall, we find that structural characteristics of the community as a whole predict the pace of technological development: In particular, the most productive communities are large ($\beta(\text{SE}) = 0.24(0.005)$, $z = 50.9$) and diverse ($\beta(\text{SE}) = -0.12(0.003)$, $z = -35.2$). By contrast, individual community members' expertise has a modest effect on technological development ($\beta(\text{SE}) = 0.008(0.004)$, $z = 2.0$). In further analyses, we find that how much an individual contributes to the community is best predicted not by their degree of expertise, but rather by the fit between their expertise and the rest of the community's activities. Our work tests theories of cultural evolution and economic complexity at scale and provides a methodological basis to study the interplay between individual expertise and community structures.

S2.4 - Learning how to behave: cognitive learning processes underlying adaptation to social norms and intergroup bias

Uri Hertz¹

¹University of Haifa

Social norms provide groups of individuals with behavioral prescriptions, and govern the way individuals perceive and interact with groups. Changes to social settings caused by migration or cultural change force us to adapt to new social norms, and to interact with individuals whose group identity is different from our own. I will discuss two studies looking at the way cognitive learning processes account for the way people learn about behavior associated with groups, such as social norms, and the way features of the norm's prescribed behavior affect learning. Using a multiplayer game, where bot-players displayed a variety of social norms, and computational learning models, we found that active behaviors were learned faster than omissions, and harmful behaviors were more readily attributed to all group members than beneficial behaviors. In a second study, we examine whether different learning mechanisms are used for group-level attribution for in-group players and out-group players. Our results and approach provide a cognitive foundation for learning and adaptation to descriptive norms and can inform future investigations of group-level learning and cross-cultural adaptation.